Community Petascale for Accelerator Science and Simulation (COMPASS) Project

Year One March 2008 Progress Report

For Office of Science (SC) of the U.S. Department of Energy
Funded by Offices of Advanced Scientific Computing Research (ASCR), Basic Energy Science (BES) High Energy Physics (HEP), and Nuclear Physics (NP)

Effective Date: March 1, 2008

Revision 1.0

1 Objectives

The objective of this report is to provide a comprehensive summary of the ComPASS project progress and the methodology to track it. A summary of the progress and related issues are described in section 2. An overview of the project management methodology is given in the section 3. Details of the project performance as reported by the principal investigators of collaborating laboratories are given in section 4. The list of related publications and presentations are given in section 5.

2 Executive Summary

2.1 Overview

The ComPASS project aims to advance accelerator computational capabilities from the terascale to the petascale to support DOE accelerator operation and design priorities for the next decade and beyond. In order to achieve this goal, the ComPASS program focuses on developing a comprehensive set of interoperable components for beam dynamics, electromagnetics, electron cooling, and advanced accelerator modeling. The components required for building this simulation suite include:

a) Development of petascale capable individual physics process modules (either new, or improve the performance of existing, from SciDAC1, modules). This step requires the utilization (development) of petascale capable algorithms and solvers.

b) Development of interfaces and framework code that will allow the operation of the physics codes above in the same environment (prototypes of such environment and interfaces were developed under the SciDAC1 project)

c) Development of physics applications with complexity that increases as the multi-physics, multi-scale modeling capabilities are developed.

Due to funding limitations, the development of comprehensive documentation and user interfaces has been descoped.

In the first six months of the project, we focused on improving the performance of our single-physics modules, began the design of common interface development for framework components, and applied our codes to single-physics model applications and began multi-physics application model development. We also developed a comprehensive project WBS and tracked project progress through regular meetings of the L2 managers and the management team of the project.

2.2 Activity report – milestones reached

There are three ComPASS high level activities:

1) Electromagnetics
2) Beam Dynamics
3) Advanced Accelerators

Each of the above activity groups includes both development and application activities and deliverables. The first two activity groups receive support from all four offices funding ComPASS

(ASCR, BES, HEP, and NP), while the third receives support only from HEP. A list of milestones and deliverables reached in the first six months together with a short summary follows. The detailed description of the activities leading to each milestone can be found in the individual research group reports presented in section 15.

2.2.1 Electromagnetics

This section presents the electromagnetics activity group accomplishments and milestones reached in the first 6 months of the project. Note that ILC related calculations are immediately applicable to any SCRF accelerator R&D effort involving similar rf cavity designs. In addition, the capabilities developed will be beneficial to any electromagnetics application problem that involves complex geometries and a large number of degrees of freedom.

1) Completed first ever calculation of trapped modes in an ILC cryomodule (HEP deliverable, SLAC lead). This became possible because of the development of new, efficient parallel linear solvers in the frequency domain for the finite element code Omega3P.
2) Impedance calculation of the ILC damping rings (HEP deliverable, SLAC lead): obtained the impedance budget for vacuum chamber components. The second component of this work involves the determination of the pseudo-Green's function wakefield used as input to beam stability studies. Such calculations require very short bunch lengths, thus necessitating the development of a moving window algorithm; this work will be completed in subsequent phases of the project's development.
3) Code benchmarking with experimental data: developed Omega3P/S3P model of a JLab 7-cell high gradient fully dressed SRF cavity and compared with experimental measurements. The results were in good agreement. (NP funded, but generic task, SLAC and TJNAF lead).
4) Completed accurate calculation (part 10^5) frequency computation of crab cavities using a time domain, finite difference code (VORPAL). The calculation improved on previous low-fidelity calculations and opened the opportunity for data comparisons (HEP deliverable, Tech-X lead).
5) Developed high-resolution simulations of the Facility for Rare Isotope Beams (FRIB) low beta cavities and the FRIB RFQ (radiofrequency quadruplole) and utilized them to assist design (NP deliverable, ANL lead).
6) Multipacting studies of SCRF cavities:
    a. Developed ILC applications using finite element approach (Track3P) (HEP deliverable,.SLAC lead)
    b. Modeled JLab cavity using self-consistent, time-domain, finite difference approach (VORPAL) (NP deliverable, TJNAF and Tech-X lead).
    c. Completed HOM coupler modeling in the SNS superconducting cavity, using finite element approach (Omega3P, Track3P). The results helped explain the behavior of the cavities during commissioning of the linac (BES deliverable, SLAC lead).
7) Developed optimization software for cavity shape determination (ASCR deliverable, SLAC and LBNL –TOPS- lead).
8) Completed porting of ComPASS codes to leadership high-performanc-computing platforms, and performance optimization on these platforms (see Table 1) (HEP funded but generic task, SLAC and Tech-X activity).

2.2.2 Beam Dynamics

This section presents the beam dynamics activity group accomplishments and milestones reached in the first 6 months of the project.

1) Developed a multi-physics application for the FNAL Tevatron, including beam-beam effects, impedance, and realistic lattice and beam helix descriptions (from data measurements). Performed multi-bunch simulations (up to 36 on 36 bunches) and successfully model measurements of bunch emittance growth (HEP deliverable, FNAL lead; collaboration with FNAL Tevarton researchers). Note that the first year ComPASS deliverable called for the completion of single bunch multi-physics simulations. We are ahead of schedule because of ILC application de-scoping which resulted in re-direction of resources, thus enabling the necessary development for this application.
2) Electron cloud capability development and applications (HEP deliverable, LBNL lead, Tech-X , FNAL, and UCLA effort)
    a. development and ILC and LHC applications:
        i. Begun development of applications for the ILC Damping Ring. This effort was interrupted, and resources were re-directed to other applications and algorithm development.
        ii. Completed study head-tail instability using a simplified model of the LHC at injection, using the fast co-moving frame model (see below).
    b. development of new algorithms
        i. completed new second order integrator accurate in the relativistic regime
        ii. introduced the utilization of a co-moving frame in the electrostatic approximation to speed up calculation of the effect
    c. Begun comparisons and benchmarking of electron cloud modeling packages. Good agreement shown in most cases.
3) Beam-beam modeling applications:
    a. Developed model for RHIC, including nonlinear optics and beam-beam interactions among 3 coupled bunches; studied and optimized tune working point (NP deliverable, LBNL and BNL lead)
    b. Developed ELIC single-IP beam-beam modeling application. Studied beam parameter regimes where coherent instability becomes important (NP deliverable, TJNAF lead utilizing LBNL support)
4) Space charge modeling applications: completed studies of space-charge effects in the ILC Ring to Main Linac (RTML) line, to determine the effects on the dispersion matching section. (HEP deliverable, FNAL lead).
5) Multi-physics application development for LHC. Developed a new model of intrabeam scattering in beam-beam simulations targeting LHC beam regime applications (HEP deliverable, SLAC lead).
6) Electron cooling code development and applications (NP deliverable, but mixed funding from HEP and NP, Tech-X lead, with BNL participation):
    a. Completed prototype electrostatic PIC module in VORPAL and benchmarked performance
    b. Begun to develop and study models of the modulator section of the coherent electron cooling concept.
    c. Studied the effect of undulator fields in e-cooling
    d. Studied the effects of finite interaction time in Coulomb collisions

7) Framework development: completed the refactoring of the Synergia2 framework and begun component interface development, with the first prototype components coming from a Synergia2 solver and a ML/Impact rf map generation module. (HEP and ASCR funded but generic task, FNAL, Tech-X, and ANL lead).
8) Poisson solver development.
    a. Completed a new FFTW based solver for open or closed boundary conditions in the Synergia2 framework. Scalability up to 1000 processors (HEP deliverable, FNAL lead).
    b. Development of new fast Poisson solver for beam-dynamics frameworks: begun by benchmarking the FFT-based solvers in existing codes (such us BeamBeam3D and IMPACT), as well as multigrid-based solvers (ASCR deliverable, LBNL lead).
9) Completed porting of ComPASS codes to leadership high-performanc-computing platforms, and performance optimization on these platforms (see Table 1) (HEP, NP, BES funded,
9) FNAL, LBNL, ANL lead).
10) Code parallelization: begun parallelization of elegant physics and elegant SDDS (Self Describing Data Sets) file operations. The API for parallel SDDS was defined, and a prototype library was benchmarked on a variety of file systems (NFS, PVFS, and GPFS). (BES deliverable, ANL lead)

2.2.3 Advanced Accelerators

The section presents the advanced accelerator activity group accomplishments and milestones reached in the first 6 months of the project.

1) Code benchmarking: we completed comparisons of the major plasma physics simulation codes (both full PIC and quasistatic) using a 3D laser wakefield acceleration problem both in the linear and non-linear regimes. The study resulted in better understanding of the specifics of problem description for different codes, and produced good agreement in the physics predictions (HEP deliverable, UCLA, TechX, and LBNL lead.)
2) Developed a new model of how to load electrons on the non-linear wakefield generated by the drive beam. Applications of this model resulted in high efficiency and small energy spread in the trailing beam. (HEP deliverable, UCLA lead).
3) We performed high-resolution simulations to study relevant physics in plasma experiments:
    a. Laser wakefield accelerator (LWFA) experiments: studied the effects of gas jet density in momentum spread and the performance of new injection mechanisms (Tech-X, UCLA, LBNL lead, in collaboration with non-ComPASS funded researchers at SLAC, LBNL, and ANL).
    b. Plasma afterburners. Modeled experiment design based on the algorithm described in (2) above (HEP deliverable, UCLA lead).
4) Developed and implemented pipelining algorithms for plasma wakefield and laser wakefield acceleration. The algorithm was implemented into the beam-driver of both the basic and quasi-static versions of QuickPIC. The results were benchmarked against the non-pipelined version of the code, and the performance scaling studied. (HEP deliverable, UCLA lead).
5) Code porting and optimization. All major codes were ported and benchmarked on a variety of leadership class high-performance-computing facilities (see Table 1). (HEP deliverable, Tech-X, UCLA lead).

2.3 ComPASS tools

A list of all codes and frameworks developed or maintained under ComPASS is listed below. Note that some family of codes share components (the naming convention will make the relations obvious) but do not constitute frameworks, while others are true frameworks. One of the overall project deliverables is to consolidate the functionality of individual codes into general frameworks.

Table 1. A table of ComPASS codes and frameworks

Code	**Description**	**Platforms utilized**	**Users**	**Developers**
Impact-T	Beam dynamics using time as the independent variable. Applications: ALS streak camera, APS upgrade, LCLS injector, Fermilab A0 experiment, and BNL, UCLA, Cornell, UW-Madison FEL, PSI FEL, SPARC/X , Fermi/Eletrra photoinjectors.	Desktop PC MPI clusters and single processor PCs.	LBNL, ANL, BNL, SLAC/LCLS, NIU/Fermilab, LANL, UCLA, Cornell, UW-Madison, PSI, INFN-Frascati, Fermi/Elettra	LBNL
Impact	Beam dynamics simulation in RF linac using position as the independent variable. Applications: Berkeley FEL studies, LEDA Halo experiment, ILC, SNS, RIKEN cyclotron injector, and RIA driver, JPARC, GSI, CERN SPL, CSNS linacs	Desktop PC MPI clusters and single processor PCs.	LBNL, Fermilab, NIU, ORNL, MSU, KEK, RIKEN, GSI, CERN, IHEP	LBNL
ML/Impact	Beam dynamics problems involving high-order optics and space charge, realistic magnetic field profiles, realistic cavity simulations, fitting and optimization. Application: RHIC	Cray XT4, IBM SP4, Blue Gene, Mac.	LBNL, U.Md, LANL, Tech-X	LBNL, Tech-X, U.Md

	upgrade.			
BeamBeam3D	Beam dynamics in hadron colliders, beam-beam and impedance effects. Applications: Tevatron, KEKB, RHIC, LHC.	Desktop MPI PC clusters, Cray XT4, IBM SP3, SP4, and BG/L.	FNAL, JLab, LBNL	LBNL, FNAL
PLIBB/ Nimzovich	General-purpose parallel tracking framework for beam-beam simulation. Application: RHIC wire compensation beam-beam experiments	NERSC supercomputers, Linux	SLAC	SLAC
Synergia	Synergia2 is an extensible framework for beam dynamics. It is driven by Python, allowing for arbitrarily complex simulations. The framework includes non-linear single-particle physics, beam generation and diagnostics, and collective effects modules. Single-particle optics is provided by CHEF, which includes arbitrary-order maps. Collective effects modules include extensive support for space-charge calculations, for open and closed boundary conditions. These calculations are available both through Synergia modules and through an interface to IMPACT. Applications: FNAL Booster, ILC DR and RTML, FNAL A0 photoinjector.	The framework is optimized for massively parallel computations utilizing MPI. It is currently ported on desktop PC MPI clusters and Unix supercomputers supporting shared libraries (IBM SP3 and SP4).	FNAL, NIU, IIT, ANL.	FNAL, Tech-X, ANL.

Elegant	Beam dynamics with focus on electron beams. Applications: FEL driver linac and energy recovery linac design.	The parallel version of elegant is developed under ComPASS	ANL, FNAL, Cornell, JLAB, BNL, Daresbury, JAERI, SLAC, DESY	ANL
Vorpal	VORPAL is a framework that can be configured in different ways at run time to solve a variety of problems. It provides models for neutral gases and plasmas interacting with electromagnetic fields and with each other through collisions. Applications: Laser Wake Field Accelerators (LWFA), EM cavity calculations (ILC, JLab), Electron cooling (BNL), Electron gun modeling, Multipactoring (JLab), Analysis of ecloud measurements.	VORPAL runs on laptops to supercomputers, including OS X, AIX, Linux operating systems. Particular supercomputers include Bassi, Franklin, Jacquard. VORPAL ports rapidly to other platforms. It has been ported to BlueGene.	LBNL, JLab, BNL, ANL, Tech-X	Tech-X, University of Colorado.
Upic	A flexible Framework for rapid construction of parallel pic applications. It supports electrostatic, darwin, and EMsolvers. Applications: quickpic is based entirely on the upic framework, IMPACT.	UPIC is strictly Fortran90 compliant, with no dependencies other than MPI and pthreads, so it compiles on any platform. It has been tested with at least 10 compilers (it works on all except gfortran); works on all platforms mentioned under quicpic, plus the	UCLA, LBNL, U. Alaska, Virginia Tech	UCLA

		SGI Altix and the asci Q machine.		
Osiris	Fully explicit pic code with ionization and coulomb collision packages. Applications: lwfa and pwfa modeling	IBM using the xlf compiler, AMDwith g95, pgf, and pathscale compilers. It runs on all the nersc machines at lbnl, the atlas machine at llnl, and on os-10 machines.	UCLA, USC, SLAC, IST, Michigan, Imperial College, Rutherford National Lab, and Rochester.	UCLA, IST (portugal), USC.
Quicpic	Quasi-static pic code based on the upic framework. Applications: lwfa, pwfa, and e-cloud modeling (LHC, ILC, FNAL-MI).	nersc machines and the UCLA os-10 based cluster. Since it is based on upic, the same portability statements hold.	UCLA, USC, SLAC, U.MD, IST.	UCLA, U. MD.
Omega3P	Frequency domain eigensolver for cavity mode and damping calculations. Applications: ILC cryomodule, LHC crab cavity and collimator, X-band high gradient structure, JLab 12-GeV upgrade superconducting cavity, BNL RHIC cavity, SNS superconducting cavity	NERSC and NCCS supercomputers, Linux, Windows	SLAC, ANL, BNL, TJNAF, FNAL, Cornell, University TAM	SLAC
S3P	Frequency domain S-parameter calculations. Applications: ILC superconducting cavity coupler.	NERSC supercomputers, Linux, Windows	SLAC, TJNAF	SLAC
Track3P	Particle tracking for simulation of dark current and multipacting. Applications: Coupler of ILC superconducting cavity, SNS	NERSC and NCCS supercomputers, Linux	SLAC, BNL	SLAC

	superconducting cavity, BNL RHIC cavity			
T3P	Time domain solver for wakefield computations with beam excitation. Applications: ILC cryomodule and damping ring, CLIC PETS structure, MIT photonic band gap structure	NERSC and NCCS supercomputers, Linux	SLAC	SLAC
Gun3P	Space-charge trajectory code for beam formation and beam transport. Applications: ILC sheet beam klystron gun	NERSC supercomputers, Linux	SLAC	SLAC
Pic3P	Self-consistent particle-in-cell code for rf gun and klystron simulations. Applications: LCLS rf gun	NERSC supercomputers, Linux	SLAC	SLAC
TEM3P	Integrated electromagnetic, Thermal and mechanical analysis for cavity design. Framework uses solvers from SLAC codes above. Applications: LCLS rf gun.	NERSC supercomputers, Linux	SLAC	SLAC
V3D	Visualization of meshes, field and particles for unstructured meshes. Applications: all SLAC codes.	Linux.	SLAC, BNL	SLAC

2.4 Budget summary

The HEP budget and personnel support for the three different high level activities for FY07 is shown in Figure 1. ComPASS FY07 HEP budget summary. The NP budget for FY07 is $179k, supporting 1.2 postdocs and 0.1 research scientist, equally split between electromagnetic and beam dynamics activities. The BES budget is $100k supporting 0.5 postdocs (beam dynamics) and 0.2 research scientists (electromagnetics).

DOE Support FY 2007 (Actual)				
Personnel Support from DOE Grant				
(in FTEs)				
Activity:	Electro - magnetics	Beam Dynamics	Advanced Accelerator	**TOTAL**
Faculty (tenured or tenure-track)	0.00	0.00	0.00	**0.0**
Research Scientist	3.32	2.56	0.90	**6.8**
Postdocs	0.00	0.00	1.00	**1.0**
Graduate Students	0.00	0.00	1.00	**1.0**
Engineer / Computing Professiona	0.00	0.40	0.00	**0.4**
Administrative / Technician	0.00	0.00	0.00	**0.0**
Other	0.00	0.00	0.00	**0.0**
DOE/HEP Funding (per activity):				
SWF(Actual, in k$)	688	913	306	**1,907**
M&S (Actual, in k$)	0	12	6	**18**
Travel (actual, in k$)	22	22	8	**52**
TOTAL	**710**	**947**	**320**	**1,977**

Figure 1. ComPASS FY07 HEP budget summary

3 Project Management

3.1 General Project Information

The ComPASS project officially began on September 1, 2007. By this date, eight institutions within the collaboration received funding for the year one. This funding level is substantially reduced from the original funding level assumed in the original proposal. Due to lack of funding, several tasks in the original proposal are descoped or modified. Also, several interdependent tasks were either modified or moved to maximize the project output.

The project Work Breakdown Structure (WBS) is organized according to the deliverables by the institutions and further broken down by the four funding types HEP, NP, BES, and ASCR. These tasks are tracked in the WBS. The executive team decided to keep track of the de-scoped tasks, so that the milestones can be properly tracked. Each descoped task is marked as a task with duration of two days and is considered complete. These tasks may be activated in the future if additional funding becomes available. The project work breakdown structure (WBS) is used to track the milestones quantitatively.

Based on the arrival of funding, each project year begins on September 1 of each calendar year and ends on August 31 on the next calendar year.

The overall management of the project is done at Fermi National Accelerator Laboratory (FNAL) which is the lead institution. Coordinating Principal Investigator and DOE Contact for the project is Dr. Panagiotis Spentzouris. He is assisted by Dr. Bakul Banerjee and FNAL financial office staff members as necessary. The project management documents are maintained online. The virtual office where the project management duties are generally performed collectively is referred to as the Project Office.

3.2 Financial Status Summary

The ComPASS project funds are delivered directly to the collaborating institutions. Each Institutional Principal Investigator (PI) is responsible for managing the allocated funds to appropriate tasks.

Each collaborating institution uses its official accounting system as a basis for collecting cost and effort reporting data. Each PI, with the help of the institutional financial officer establishes a relationship between the WBS elements and the institutional account code under its accounting system. Periodically, PIs submit the fully burdened total cost incurred to accomplish the tasks to the ComPASS project office. The fully burdened cost implies not only what is paid to the employee or the contractor, but also all associated direct and indirect costs. The project office summarizes the data and prepares financial reports.

Periodically, each institution should also report the effort spent by individuals working on various tasks. The effort reporting should be done as a percent of equivalent Full Time Employee (FTE) Year. Please use the cumulative effort spent by each individual over the period of reporting during each fiscal year.

[Add the summary spreadsheet of funds available and the actual expenditure]

3.3 Work Breakdown Structure(WBS) and Task Completion Status

The ComPASS WBS is based on the milestones and tasks established in the original proposal. ComPASS for the year one of the project.

For the purpose of this project, each task in the WBS is categorized as:

1. In scope: The funding is available and work is progressing as planned
2. Descoped: The funding is not available and the task will not be done. These tasks are given a during of 2 days and marked as 100% complete
3. Modified scope: A task has a modified scope, if only a partial funding is available and the scope may be changed or work may be accomplished under the umbrella of another funding.

A summary of modifications is provided for each section. It should be noted that it is not feasible to maintain detailed records of the changes described above. The ComPASS Executive Committee has approved the initial ComPASS WBS.

3.4 Project Issues

4 Accomplishments – Year 1

4.1 ANL

4.1.1 Progress and Results

4.1.1.1 Electromagnetics – NP

In the framework of this project, we have studied resonators for future hadron and heavy-ion accelerators. There is strong demand for medium energy proton, H-minus, and ion accelerators for various applications. Typical examples are the FRIB (Facility for Rare Isotope Beams) accelerators, the 8-GeV Proton Driver at FNAL as a neutrino factory and injector to the synchrotron and the upgrade of the SNS linac. The majority of superconducting structures being developed for medium-energy accelerators are based on quarter-wave and half-wave cavities with TEM-like transmission line modes. These structures are much more efficient in the low-velocity (beta<0.7) region than elliptical resonators despite the more complex mechanical structure of the cavities. To cover the velocity range from β_G=0.02 to β_G =0.7, several different types of resonators are required. In addition, the operating frequency varies from 50 MHz to 800 MHz, depending on the application. Cavity design also depends on the operation mode, which can be pulsed or cw. Electromagnetic and mechanical design of TEM-class cavities is more complicated than the design of elliptical cavities. The current approach in designing TEM-class cavities is a full-scale prototyping, which requires significant resources. Our particular interest is in new parallel solvers for thermal and mechanical analysis to couple the existing electromagnetic codes to form a self-contained set of tools required for engineering actual prototypes. In the framework of this project we are collaborating with SLAC and STAARinc company on developing highly scalable versions of Omega3P and ANALYST codes.

Main results of the past year are:

- We have successfully used the commercial version of Omega3P called ANALYST as well as M-W Studio to perform high-resolution calculations of the FRIB driver low-beta cavities.
- Optimized the peak surface fields with high accuracy in a 48.5 MHz cavity for FRIB post-accelerator.
- Performed electro-dynamics calculations for a Radio-Frequency Quadrupole (RFQ) using the Parallel version of ANALYST.
- The results of these studies contributed to three journal papers. Two of them are published in PRST-AB and Computer Physics Communication.

Future plans:

- Apply ANALYST for the design of TEM-class SC resonators operating at 72.75 MHz for the FRIB post-accelerator.
- Optimize the design of a 4-gap cavity for the FRIB driver operating at 57.5 MHz.
- Perform structural analysis of the FRIB post-accelerator cavities using ANSYS and other available software (S3P,..).
- Optimize the frequency tuning methods of TEM-class cavities.

4.1.1.2 Beam Dynamics - BES

WBS 1.1.3.1 Finish parallelization of CSR and longitudinal space charge elements in **elegant:** Due to the delay in receiving SciDAC funding, this proposed work was supported by another funding source. This was necessary to address the need for this feature in APS upgrade studies.

WBS 1.1.3.2 Parallelize SDDS file operations in Elegant: The plan is to develop a library for parallel SDDS file operations: SDDS ("Self Describing Data Sets") is a file protocol used by **elegant** and other APS simulation software, as well as the APS control system. The serial nature of SDDS was found to be a bottleneck for SDDS-compliant simulation programs such as parallel **elegant**, which involves a lot of SDDS IO operations. The master processor had to perform I/O, creating a bottleneck that was particularly evident when large numbers of simulation particles were needed.

A library for reading and writing SDDS files in parallel was developed. The Application Programming Interface (API) for Parallel SDDS was derived from serial SDDS with minimal changes. The API defines semantics for parallel access and is tailored for high performance. The underlying parallel IO is built on MPI-IO.

The performance of parallel SDDS was compared to parallel HDF5, one of the most widely-used formats, using ANL's Jazz cluster with PVFS version 1 file system based MPICH1 MPI-IO. Parallel SDDS shows better scalability than parallel HDF5 for row-major ordered files, and had better performance when the number of processors was greater than 20 for reading and 32 for writing. This is shown in Figure 1.

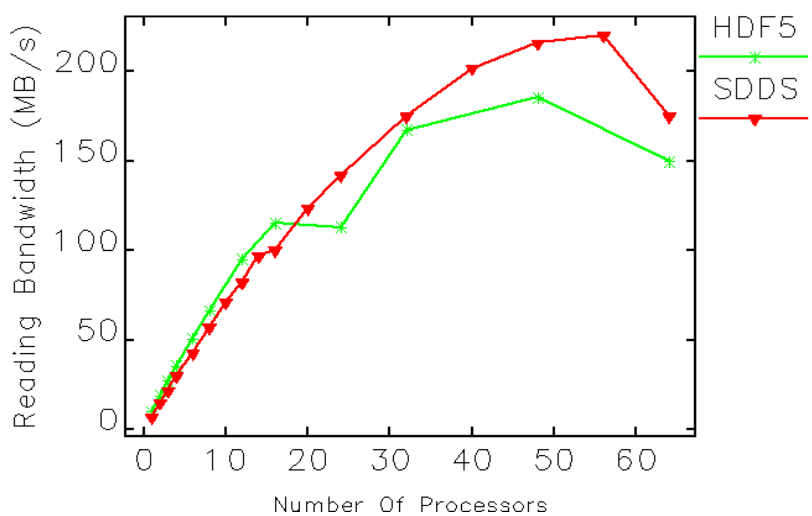

Figure 2: Read performance of parallel SDDS compared to parallel HDF5 on ANL's Jazz cluster.

The writing performance of parallel SDDS was also studied on different file systems, namely, NFS (on APS's apex cluster), PVFS (on ANL's Jazz cluster) and GPFS (on ANL's BlueGene/P cluster). It turns out that NFS scales very poorly, PVFS shows much better parallel performance, but GPFS shows the best performance on BlueGene/P. The throughput of parallel SDDS grows linearly with the number of processors for writing a 1.2GB file when the number of processors is less than 32.

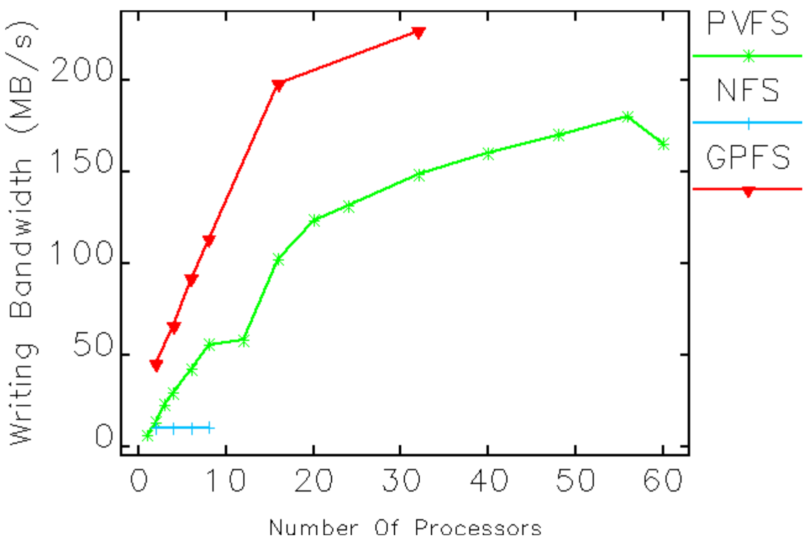

Figure 3: Write performance of parallel SDDS for different file systems.

However, the reading performance of parallel SDDS can still be improved. For example, currently all the processors read the SDDS layout, parameters, number of rows and arrays simutanously. This can be improved if only one process reads this information and broadcasts them to all others, so the IO overhead will be greatly reduced. Another improvement would be reading/writing SDDS in column order so that many internal data conversion is no longer necessary; this will be done in the SDDS-3 libraries, which are currently in development.

The integration of SDDS parallel I/O with parallel **elegant** (**Pelegant**) is accomplished as follows. Parallelization of **elegant** uses the master/slaves model. The gradual parallelization strategy, essentially an element-type-by-element-type strategy, proved very beneficial in getting parallel features into the hands of users quickly and had significant payoffs for the APS upgrade and other investigations. As a result of this approach, the tracking routine in **Pelegant** keeps track of whether the particles reside on the master or slaves. By looking at the next element classification, it can determine whether to scatter the particles, gather the particles, or leave them where they are.

Without parallel SDDS IO, all particles have to be read by the master and gathered to the master to dump the results. The limitation of this strategy is not only the communication overhead between processors, but also the memory bottleneck on the master for a simulation with a very large number of particles. This particularly restricts us when a supercomputer with limited memory on each node is used (e.g., a 32 bit architecture).

The adaption of the new parallel SDDS library resolved these issues. The efficiency was improved by the elimination of communication overhead for scattering and gathering particles, and scalable throughput was gained through parallel I/O operations. The most important advantage is that the requirement for the master to hold all the particles in memory becomes unnecessary, which makes it possible to run simulations with billions of particles. Another benefit of using parallel I/O operations is that the user of the software will have simulation result in the same format as in the serial version, i. e., the user does not need to combine several files to get final simulation result.

Considerable efforts were made to allow **Pelegant** work with parallel SDDS I/O**.** In a simulation with a very large number of particles, since no central process could hold all the particle information because of the memory requirement, we had to parallelize all the remaining serial elements used in the simulation. (In this case, simulation of FEL driver linacs such as LCLS.) Besides the final output of the simulation, some intermediate output from elements had to be adapted to use parallel I/O, which was done by the master before and it becomes impossible for large simulations. The burden of calculating the statistics information is also distributed to all the working processors instead having the master do all the work (in **elegant,** the I/O operations and diagnostic elements are classified as the same type of elements, so this becomes possible after the I/O is parallelized). For simulations with billions of particles, this will save both communication time and calculation time significantly compared with the version without parallel I/O.

1.1.3.3. Study scaling and optimize code using LCLS as benchmark problem

Performance and applications of **Pelegant** with parallel SDDS I/O is studied as a part of this task. Due to the enhanced computing capability of parallel **elegant**, simulations with a large number of particles is now possible. The necessity of running simulations with millions of particles was confirmed in the modeling of the microbunching instability for FERMI@ELETTRA, performed on the APEX cluster. Note that the studies started with the FERMI@ELECTRA linac because of delays in obtaining the elegant latitude for LCLS. We increased the number of particles that can be run on our cluster from 5 million to 60 million and found the noise is reduced significantly. However, even on a 64-bit machine with 16GB memory on each node, **Pelegant** without parallel SDDS can only simulate with 60 million particles because of the memory bottleneck on the master. In addition, the startup and output phases of the simulation are slow due to page swapping.

For these simulations, we found that 20M particles were needed to reduce noise and limit instability saturation, allowing us to obtain reliable instability gain curves down to a 25 micron modulation wavelength. The greatest value of increased particle number is not noise reduction, but rather the ability to reduce bin size and thus modulation wavelength for gain studies. Also, CSR gain increases as wavelength decreases, but rolls off "at some point". Using Pelegant with parallel SDDS IO, we have already run 400M particles and anticipate being able to run 1 billion particles. This should allow us to perform microbunching gain scans below 1 micron modulation wavelength, which is approximately the wavelength of the laser in the LCLS laser/undulator beam heater.

We did a preliminary scalability test with 100 CPUs on our Apex cluster at APS. The results are excellent, as shown in Figure 3. This test used the NFS file system. Faster I/O throughput is expected with more advanced file systems. A thorough performance test will be conducted on a supercomputer in the near future.

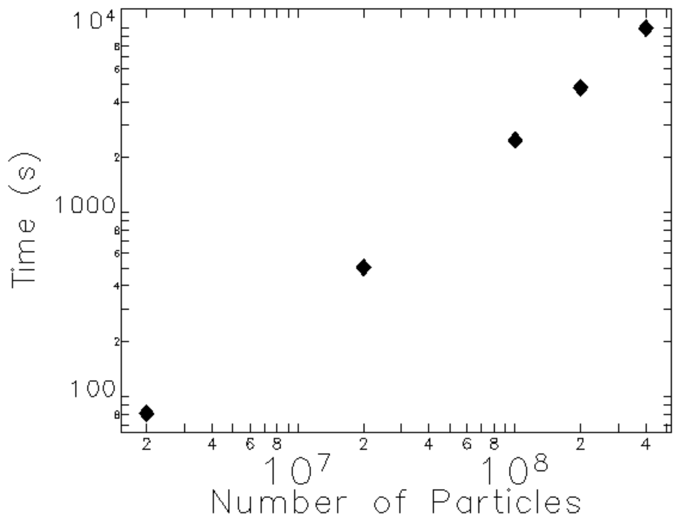

Figure 4: Scalability test of Pelegant with parallel IO for FERMI@ELETTRA, showing the wall clock time vs the number of simulation particles

Proposed work for the remaining for this funding year

Our primary on-going work for the remainder of the funding year is studying the scalability and optimizing the code using LCLS as benchmark problem. All the elements required by the LCLS simulation have been parallelized and the validation with the serial version is on-going. Part of this work will be changes to allow generating identical random number sequences in serial and parallel runs. We then would port the code to the Franklin supercomputer at NERSC to run simulations with a very large number of particles. Our goal is to simulate the microbunching instability to very short wavelengths, ideally under 1 micron. We hope to shed light on an unexpected phenomenon observed during LCLS commissioning, which appears to be related to microbunching at optical wavelengths.

4.1.1.3 SAP - ASCR

Background: A key concept in the COMPASS project, which focuses on developing a comprehensive set of interoperable components for end-to-end multi-physics simulations involving beam dynamics, electromagnetics, electron cooling, and advanced accelerator modeling, is the notion of software infrastructure for multiphysics-based accelerator modeling on petascale architectures. This effort demands a high degree of team collaboration in order to manage the integration of a wide range of accelerator physics applications, frameworks, and numerical libraries. An important facet of this work is devising common interfaces to encapsulate pre-existing physics modules, thereby facilitating code reuse among the computational accelerator team and fostering the incorporation of new capabilities developed over the lifetime of this project.

Such common interfaces will then help accelerator physicists to explore performance tradeoffs among different algorithms and implementations for various simulation scenarios and target petascale machines. For example, the parallel solution of large-scale collective effects calculations constitutes a significant fraction (typically 50-90 percent) of overall simulation time for Synergia, which is responsible for beam dynamics modeling in the COMPASS project. It is not clear a priori how to select and parameterize Synergia's current collective effects solver, an FFT-based Poisson solver, optimally for a given simulation instance. Moreover, efficiently calculating collective effects will benefit from considering new algorithmic ideas as well as others, such as exploiting periodicity that arises during repetitive solves. Overall simulation time could also be reduced by performance optimizations to exploit features of massively-parallel systems such as the IBM BG/L and Cray XT3.

Approach: The approach builds upon ongoing work on Synergia2 and MaryLie/IMPACT and will extend to other subdomains of the COMPASS project. The work employs the Common Component Architecture (CCA) to develop infrastructure for component-based accelerator simulations, with a goal of facilitating the interaction among multiple physics modules, such as space charge, electron cloud, and wakefield effects, and providing easy access to computational tools under development by SciDAC2 projects. For example, through common interfaces, prototype TOPS solver components offer easy access to a range of (non)linear solvers developed by different groups at different institutions, including hypre, PETSc, and SuperLU.

Technical Progress: Initial work has focused on defining interfaces and initial high-level componentization of the Synergia beam dynamics application (in collaboration with TASCS). Highlights of newly developed components are: F90-based beam optics components (quadrupoles and drifts) built on the MaryLie/Impact application (LBNL), C++ and F90 particle store components based on the Synergia2 framework (FNAL), and a newly implemented C++-based space charge solver that makes use of Synergia2, PETSc (ANL), and FFTW. Component interfaces enabled us to easily encapsulate and share desired functionalities from the existing codes. Further details are available in references [1,2,3,4] below.

A prototyping effort that will eventually lead to the development of reusable electron cloud physics components began during this reporting period. Ionization routines from the TxPhysics library were wrapped to allow their use as function mapped components.

Future Plans: Future work will focus on refactoring existing component interfaces for Synergia FODO-cell prototype and developing medium-grain components for particle tracking. We will also continue development of beam optics and electron cloud components and will complete the initial parallel solver component implementation. We will evaluate the performance of these new accelerator components, and we will use these results to motivate further development (in collaboration with TASCS) of computational quality of service infrastructure to facilitate the dynamic composition, reconfiguration, and substitution of accelerator components during long-running simulations.

4.1.2 Year One Scope Modifications - ANL

4.2 BNL

4.2.1 Progress and Results

Electron gun development: The RHIC electron cooler optics was optimized for different bunch charges (3 nC, 5nC, 7nC and 10 nC), bunch shapes (beer can and ellipsoid) and cathode temperatures (0.1 eV and 0.3 eV). The results are published in the feasibility document. In addition the error sensitivity was explored. An ILC electron injector was designed and optimized. A solution was found that, for the ideal machine, exceeds the current ILC requirements without damping ring. A publication is in preparation and currently reviewed by Ilan Ben-Zvi.

Electron cooling: In collaboration with Tech-X three problems were investigated in simulations: the effect of finite interaction times in Coulomb interactions were studied, the effect of undulators on the friction force was simulated, and the effect of space-charge compensating solenoids on the friction force was analyzed.

Beam-beam: Due to limited funding no beam-beam work at BNL was funded from COMPASS. However, beam-beam problems are investigated on an ongoing basis and we are in contact with other COMPASS beam-beam collaborators so that relevant problems are addressed.

Technical Progress in FY 2008

Electron gun development: It is planned to use the VORPAL cold on BNL's BlueGene system, and study with 3D particle-in-cell (PIC) simulations a 1 ½ cell SRF electron gun in detail. This requires roughly 500 processor hours for 2 GHz opteron chips. VORPAL has been installed and tested on a BlueGene-L system. Based on these tests, we anticipate runtimes on the 800 MHz PowerPC 450 chips of a BlueGene-P system will be 3x longer, or 1,500 processor hours. It will be important to study the effects of higher resolution, as this can be critical near the photocathode, where the dynamics of the low-energy electrons is strongly influenced by the highly nonlinear fields. For electromagnetic PIC, runtime scales as the 4^{th} power of resolution. Hence, 4 times higher resolution along each dimension will increase the nominal runtime by a factor of 256. Taking into account some loss of efficiency as the number of processors is increased to ~4000, we estimate 600,000 processor hours will be required for high-resolution runs. An allocation of 3 million processor hours is requested. This will enable extensive parameter scans at low resolution (~400 runs), with fewer at intermediate resolution (~100) runs and just 2 runs at very high resolution.

Electron cooling: In collaboration with Tech-X Corp. simulations and analysis will be done of the response of an electron beam to a moving ion in order to develop concept of Coherent Electron Cooling with application to RHIC.

Beam-beam: Due to limited funding no beam-beam work at BNL will be funded from COMPASS. However, beam-beam problems are investigated on an ongoing basis and we are in contact with other COMPASS beam-beam collaborators so that relevant problems are addressed.

For the electromagnetic and electron cooling work, in FY2008 the ERL test facility will be commissioned. The optimizer will be benchmarked with measurements from the ERL injector and further optimization will be performed. The capabilities of the new code to calculate correctly the beam emittance of the ERL will be tested. The electron cooler is expected to obtain CD2 approval. The design of the injector will be based on the optimization.

Expected Progress in FY 2009

For the beam-beam work, in FY2009 the newly developed capabilities will be used to optimize the beam-beam performance of RHIC, possibly beyond the Enhanced Luminosity parameters. We expect to operate with a total beam-beam tune spread of 0.02. RHIC is the first hadron collider, in which coherent beam-beam effects are important. Beam instabilities triggered by the beam-beam interaction were observed, and a parameter search and optimization may be needed. Beam-beam simulations will include the nonlinearities of the final focus triplet magnets. The long-range beam-beam compensator model will be further exercised.

For the electromagnetic and electron cooling work, in FY2009 the RHIC electron cooler will be optimized. The (hardware) bunch charge shaping system will be implemented and the simulations will be benchmarked with measurements. An attempt will be made to meet the requirements of the ILC injector without the use of a damping ring.

Expected Progress in FY 2010

For the beam-beam work, in FY2010 simulations for electron beam of a possible eRHIC linac option will be carried out in greater detail, with a self-consistent model for both the electron and hadron beam. The long-range compensators are studied in situations where the long-range beam-beam interactions are with non-round beams that may require wires of non-round cross section. It is also planned to simulate and beam-beam head-on compensating electron lens together with a nonlinear optics against RHIC upgrade parameters.

For the electromagnetic and electron cooling work, in FY2010 the optimization of the gun cavity geometry will be implemented. This requires a system that is able to generate a input from the optimizer variables and iterate the field solver program calculations to obtain the correct frequency, maximum field on the wall and other quality parameters.

Expected Progress in FY 2011

For the beam-beam work, in FY2011 the head-on compensating electron lens can be studied in greater detail with a self-consistent model, and with an intrabeam scattering model.

For the electromagnetic and electron cooling work, in FY2011 the optimizer will be benchmarked with measurements of the RHIC electron cooler. Optimization of the injector for the eRHIC linac will be performed.

Relationships to Other Projects

The understanding and improvement of beam-beam effects is beneficial to any colliding beam facility. We expect RHIC results to be applicable to the LHC, and eRHIC results applicable to ELIC. High-energy electron cooling is a new research topic, and progress in this area can be beneficial to other high-energy hadron facilities. Progress in electron gun development can be of interest to the ILC and linac-based light sources.

4.2.2 Year One Scope Modifications

4.3 FNAL

4.3.1 Progress and Results

The Beam Dynamics effort at Fermilab, which includes Synergia2 and BeamBeam3d development, has made progress in applications, infrastructure and module/solver development. At the beginning of the period, our application focus for Synergia2 was the ILC. Although we had planned to focus on the ILC Damping Ring, we received a request from Peter Tenenbaum for a study of the effect of space charge in the ILC Return to Main Linac (RTML.) Our study involved modifying the space charge kick scheme in Synergia2 from the original equidistant scheme to an element-by-element approach appropriate for a very complicated, but single-pass beam line like the RTML. We found that space charge effects do not interfere with the planned correction scheme in the RTML. Our results were delivered in the form of a report to Peter Tenenbaum. Since the beginning of 2008, we have shifted the focus of our application and, to the extent necessary, development activities to Fermilab's Project X.

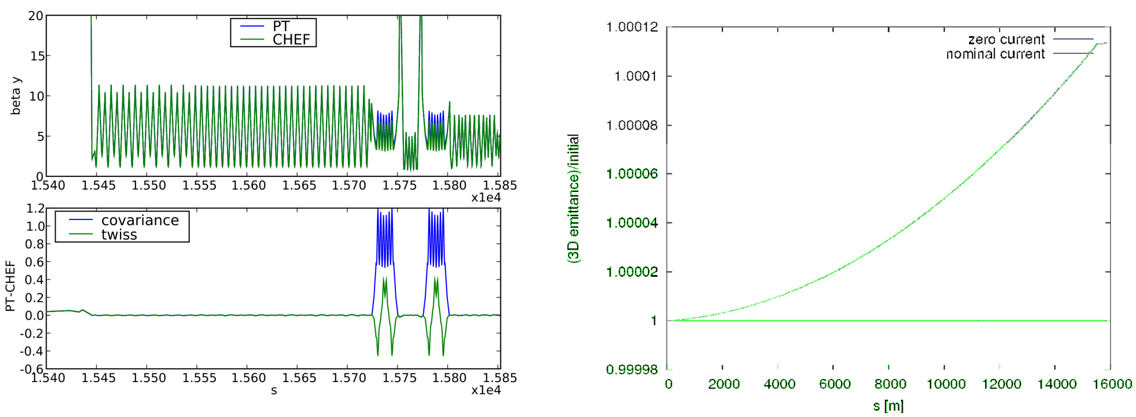

Figure 5. Study of space-charge effects on dispersion matching for the ILC RTML

BeamBeam3d application and development efforts have focused on the implementation and running of realistic six-on-six bunch simulations of the Tevatron. The simulations exhibit known features specific to the Tevatron running conditions. Preparation of a journal publication on these efforts is in progress. It contains comparisons of BeamBeam3d simulation results with data and theoretical models ("Fully 3D Multiple Beam Dynamics Processes Simulation for the Tevatron", draft available at https://compass.fnal.gov/talks-and-publications/).

In addition, we performed 36 on 36 bunches Tevatron simulations in an attempt to explain emittance growth patterns in the presence of kicker misfiring. The results are shown in Figure 7. This simulation run includes beam-beam interactions at the CDF and D0 collision points and one long-range interaction point upstream and downstream of each. The three-fold symmetry of the Tevatron filling pattern is evident as is the different behavior of the head bunch of each train: bunches 1, 13, and 25. The last bunch of each train also shows emittance growth less than the interior bunches, in

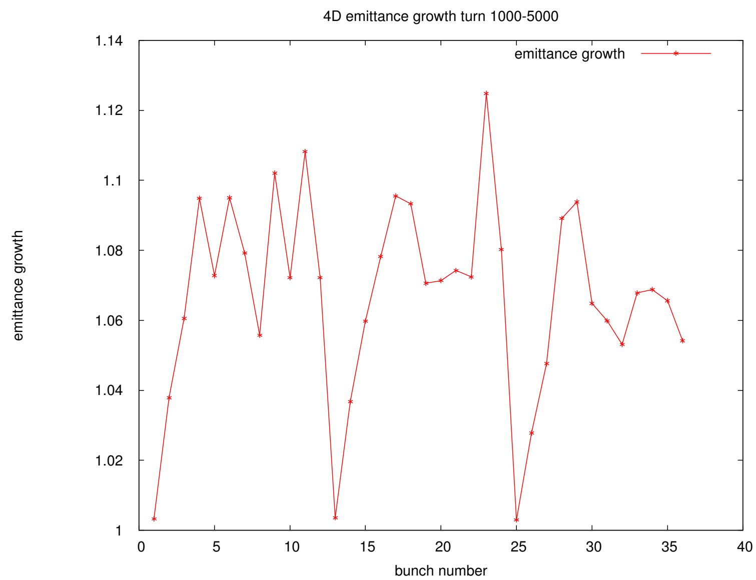

Figure 7. The 4D emittance growth as a function of the bunch number in a full 36 on 36 bunch Tevatron simulation.

agreement with observation.

The main focus of the Synergia2 infrastructure development has been the move toward

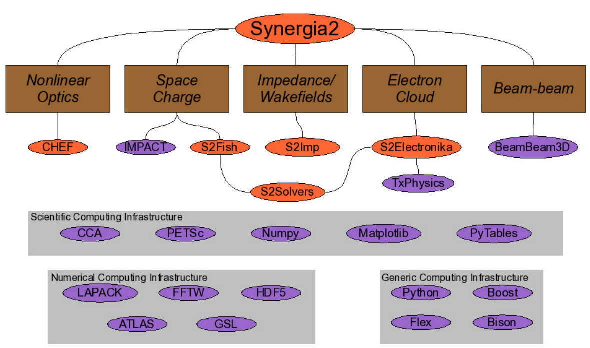

Figure 6. The Synergia framework

componentization. We have accomplished a refactorization of Synergia2 as the first major step toward a fully componentized framework. The refactored code has the additional benefits of greatly simplifying the user code needed to run Synergia2 simulations. We also include automatically documentation generation for the first time in Synergia2. As schematic of the Synergia framework is shown in Figure 6

A second task in the infrastructure area is porting. BeamBeam3d has been successfully ported to the Franklin cluster at NERSC and the BlueGene/L machine at Argonne. Eric Stern has conducted performance tests comparing these two machines with Linux clusters. He finds that the BeamBeam3d runs fastest on Franklin, and slowest on BlueGene/L, with the (relatively small) Linux cluster falling in the middle. His results for BlueGene/L are consistent with a related published study. The porting of Synergia2 is made more difficult by its dependency on several outside package and its fundamental utilization of shared libraries. We have made progress in moving Synergia2 to cross-compiling environments such as BlueGene by porting it and the CHEF libraries to the CMake build system. We are continuing to work on the shared library problem.

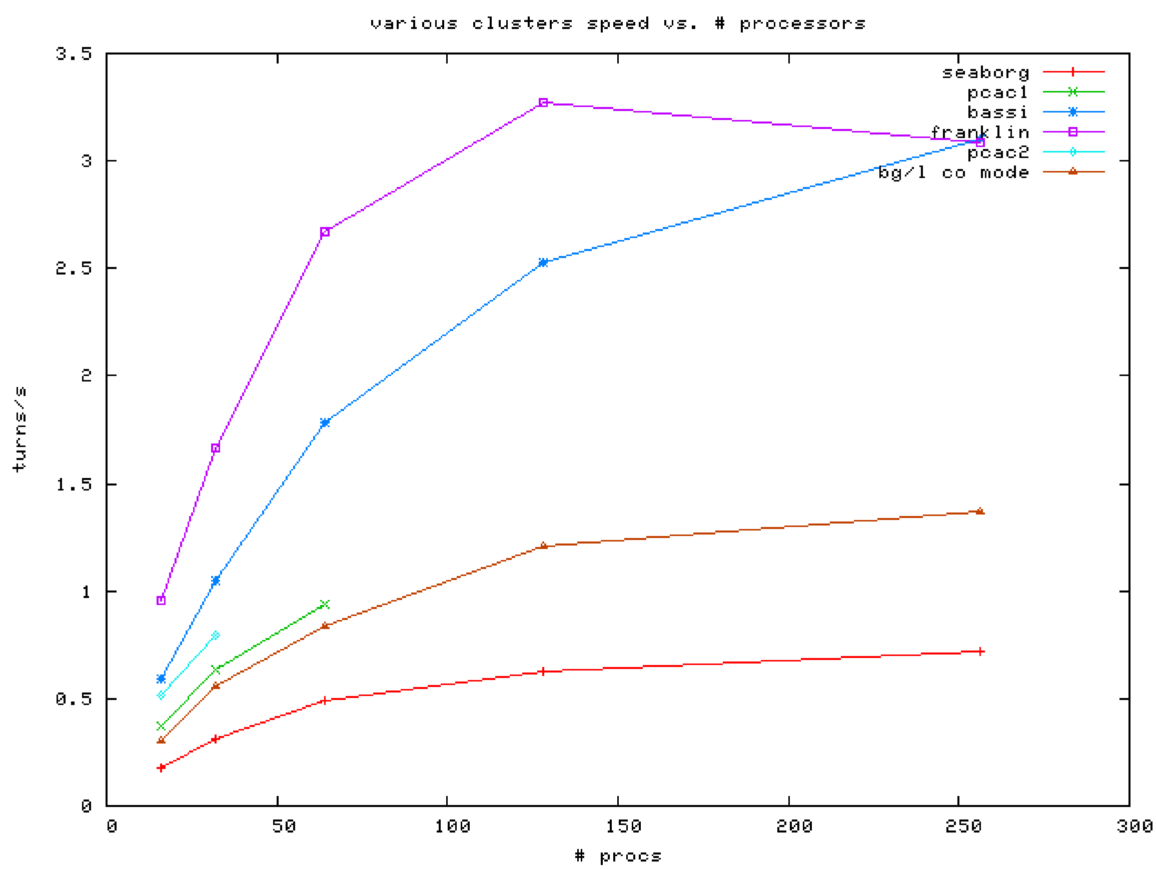

Figure 8. BeamBeam3D performance.

Our model/solver developments have included work toward new physics modules for Synergia2 and more efficient solvers for the physics modules. We have focused on solver development first because it is applicable to both space charge and electron cloud models. We have completed the development of a new Poisson FFT based solver based on parallel FFTW. The scalability is similar to our existing solvers (up to 1000 processors, depending on the platform), but the flexibility of the solver much greater. When our main application focus was the ILC, the most important issue in solver development was handling very high aspect-ratio beams. Since our application work has now shifted toward Project X, we are shifting our solver development to handling elliptical pipes, such as are found in the Fermilab Main Injector. We have a prototype of a new class of solvers that should be well suited to this type of problem. We are also optimistic that we can use similar techniques to solve the high aspect-ratio problem. We have also implemented a prototype transfer of the impedance model we previously developed for BeamBeam3d to Synergia2. Impedance effect calculations will be crucial for Project X simulations.

Fermilab also was responsible for creating and providing content for the ComPASS web portal (compass.fnal.gov), creating a baseline for the project's first year plan, organizing meetings of the management committee and the collaboration, and for creating this report.

4.3.2 Year One Scope Modifications

4.4 LANL - Descoped

4.5 LBNL

4.5.1 Progress and Results

Activities fall into 3 main areas:
(1) Beam Dynamics
(2) Electron-Cloud Modeling
(3) Laser Wakefield Accelerator Modeling

4.5.1.1 Beam Dynamics

BeamBeam3D Code: We have ported the BeamBeam3D code to a new state-of-art supercomputer, Franklin, at NERSC through the early user program. Franklin is a new Cray XT4 parallel computer with about 20,000 processors. After some modifications of the code, it runs smoothly on the computer with performance about a factor of 5 faster than on the old Seaborg machine. The nonlinear effect from the accelerator magnet element plays an important role in beam lifetime. We have added a thin lens map to model nonlinear sextupole magnet in the BeamBeam3D code.

MaryLie/IMPACT: We ported ML/I to Franklin. We did a survey of codes that use a circuit model of beam cavity interactions and started planning for implementation of a circuit model in ML/I. We enhanced ML/I's embedded envelope package to calculate not only the rms envelopes but also the depressed phase advances. We modified the beam generation command to allow scaling of the beam by the rms envelope parameters. With these enhancements we have successfully performed combined simulations involving rms matching, beam generation, and particle tracking. We also began code modifications to do away with the grid replication currently used in the space charge portion of the code. We located a 3D parallel fft package that allows for arbitrary (1D, 2D, and 3D) decomposition. Though the package did not come with a makefile for a Cray XT4, Richard Gerber of NERSC User Services examined the package and provided us with a makefile for Franklin. Using that makefile we ran the test code provided and verified that would could perform fft's on grids of up to 2048^3. We embedded the fft package into a version of ML/I and have begun implementing a version of the ML/I Poisson solver that will make use of it.

Application of BeamBeam3D to RHIC: Using the enhanced BeamBeam3D code, which includes self-consistent calculation of beam-beam interactions among three coupled bunches and linear/nonlinear external focusing, we have studied the choice of a new tune working point of proton at RHIC. Our study suggested that a working point near the integer (but with sufficient distance from the integer) would have a better beam quality (i.e. smaller emittance growth) than the current operational working point.

Modeling High Intensity/High Brightness beams: We have successfully ported the IMPACT code suite to the new Franklin machine. We have also improved the robustness of the code by fixing some defects associated with the Lorentz integrator in the code. This makes the code produce the same simulation results even under strong load imbalance situation. We have also improved our magnet quadrupole model by using a linear transfer map including the chromatic effect. This feature is important to accurately model the beam dynamics with significant energy spread in a linac or a ring.

Electron Cooling: In the area of electron cooling modeling using the Langevin approach, we have done literature studies of multi-species Coulomb collisions. We have identified the method to conserve the energy and the momentum during two species collision. We plan to carry some more numerical tests in the near future.

4.5.1.2 Electron-Cloud Modeling

We report progress into bringing the Particle-In-Cell method to handle ultra-relativistic species for the modeling of the transport of beam into high-energy accelerators. The new algorithms have been reported in a paper that has been published in the journal Physics of Plasmas [1], and are summarized below.

It was shown in [2] that for a certain class of problems involving objects (made of matter or light) propagating near or at the speed of light, the range of space and time scales spanned by the system depends strongly on the velocity of the frame of reference with regard to the system. For commonly used methods in computer physics simulations, such as for example the ubiquitous Particle-In-Cell method in plasma physics, the implication is a difference of orders of magnitude on the number of mathematical operations needed to solve a problem, based solely on the choice of the frame of reference. Since the principle of relativity implies that the laws of physics are the same regardless of the chosen frame of reference, and that the Particle-In-Cell method is based on the discretization of the fundamental laws of particle motion and electromagnetism, one might think that the solution is simple: just do the calculation in the frame which minimizes the range of space and time scales. In practice, however, we have found that the discretized equations may not preserve some fundamental properties of the continuous equations which may lead to unacceptably large errors.

We have discovered that the commonly used Boris algorithm, a second-order leapfrog integrator of the equations of motion [3], does not preserve a fundamental property of electric and magnetic field cancellation that is essential to the calculation of the orbits of relativistic species. We have derived an alternative formulation of the second-order leapfrog solver that preserves this property, and contrasted numerical results with the Boris scheme on a few simple test cases.

Assuming that waves and retardation effects are negligible, Maxwell's equations reduce to Darwin's equations [4]. However, the numerical solution of the Darwin set of equations leads to an implicit scheme which has been reported to be expensive to solve [5]. We have found a simpler system by making the following additional assumption: for each species, we have assumed that the electrostatic approximation is sufficient when the fields are computed in a co-moving frame. As a result, we have obtained a procedure for solving the fields which still retains electrostatic, magnetostatic and inductive field effects in the direction of the mean velocity of the species, is fully explicit and simpler than the full Darwin approximation.

The new algorithms were implemented in the accelerator PIC code Warp and their effectiveness were demonstrated on the modeling of the head-tail instability for a beam interacting with electron clouds in a simplified configuration of a beam in the LHC at injection.

4.5.1.3 Laser Wakefield Accelerator Modeling

SciDAC modeling of laser wakefield accelerators over the last six months has focused on modeling of plasma density ramp controlled injection to stabilize and improve beam quality, staging of injector and accelerator modules, and 10 GeV stages in support of the BELLA proposal.

Recent LWFA experiments observed electron bunches with longitudinal and transverse momentum spreads more than ten times lower than in previous experiments at central momenta of near 1 MeV/c, and supporting SciDAC simulations showed that this was possible because the plasma density down ramp at the end of the gas jet controlled the wake phase velocity and trapping threshold. Because the plasma wavelength lengthens in the downramp, controlling the ramp length and the laser focus within it allowed trapping to be precisely controlled. Detailed simulations reasonably reproduced the beam parameters and were benchmarked to THz measurements to show that the bunch is 10’s of fs long at formation, allowing use as an LWFA injector. Experiments and simulations observe high laser transmission through the gasjet, and simulations show that this allows the transmitted laser to drive a wake in a subsequent plasma to further accelerate the bunch. The momentum spread is not greatly degraded in such post-acceleration, and simulations have observed bunches with hundred-keV energy spread at energies up to ~ 20 MeV. Evaluation to higher energies is in progress. Hence, use of these bunches as an injector may greatly improve LWFA beam quality. Experiments and simulations both showed accelerator performance was also stable (three experiments over 7 days); such stability will be crucial for LWFA applications and has not been previously observed. A paper and two proceedings have been submitted on this work.
Simulations also supported development of 10 GeV LWFAs as part of the BELLA proposal for a 1 PW rep rated laser at LBNL. For the review, simulations were used to extend and detail the theoretical design for a 10 GeV LWFA. Explicit simulations at 10 GeV are forbiddingly costly due to the meter scale of the plasma needed, and hence required development of new simulation techniques. We worked with collaborators on simulations with Tech-X’s newly developed ponderomotive codes (which do not resolve the laser frequency) and Lorentz shifted codes (which use a non-stationary frame for calculation), both of which speed calculation dramatically. These simulations were used with explicit (standard Particle in Cell) simulations in 1 dimension and scaled simulations. Together, these techniques allowed multidimensional self - consistent simulation of parameters that would have been inaccessible, yielding detailed design of future experiments and laser requirements.

Simulations of self trapping and acceleration in support of 0.1-1GeV capillary and gas jet experiments also continued, with a focus on better modeling of experimental parameters and higher order beam moments.

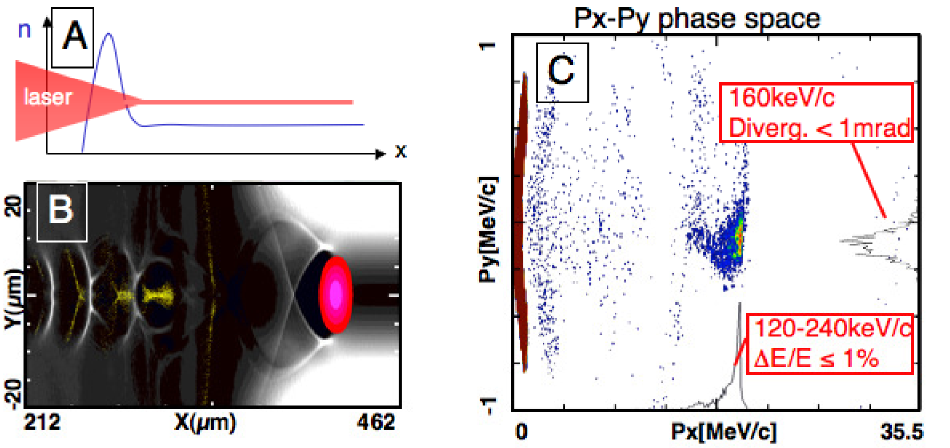

Figure 9: SciDAC simulations show that use of plasma density gradients as a LWFA injector can reduce energy spread and emittance. The down ramp is merged directly into the channel (A), and the simulation frame where the ramp merges into the channel (B) shows the laser pulse (red), plasma density wake response (grey) and particles trapped in the ramp (yellow). After several mm of acceleration in the channel, energy rises from one to 20 MeV but the narrow transverse and longitudinal momentum spreads of the bunch are retained.

4.5.1.4 SAP - ASCR

1) Nonlinear eigenvalue calculations in EM simulations: We have been pursuing literature on the solution of one-parameter nonlinear eigenvalue problems, which is relevant to the EM simulations at SLAC. One of the goals is to identify a set of algorithms that are appropriate for large sparse problems and are amendable for parallel implementations. The other goal is to detemine limitations of these algorithms and identify new research directions. The progress has been hampered by the resignation of a staff member. However, we have identified an appropriate candidate to fill the position. The candidate is finishing his PhD at UC Davis and will report to work in July. In meantime, the candidate will make himself available to meet with LBNL and SLAC staff members in order to catch up with the project.

2) Fast Poisson Solvers in beam dynamics simulations: We have begun collaborating with Rob Ryne's group in the area of fast Poisson solvers. The current plan is to benchmark the FFT-based solvers in BeamBeam3D and IMPACT, as well as multigrid-based solvers. This collaboration will eventually involve staff members from Fermi Lab. Future work will include an investigation of the need for factorization-based Poisson solvers.

4.5.2 Year One Scope Modifications

4.6 ORNL- Descoped

4.7 SLAC

4.7.1 Progress and Results

(1) Simulation of the ILC Cryomodule (1 FTE)

The determination of wakefield and higher-order-mode (HOM) effects is of utmost importance in the design of the ILC linacs. Previous calculations had been carried out for a single superconducting cavity where the spectrum and the damping factors of HOMs were obtained using the parallel finite-element eigensolver Omega3P. The calculations assumed that the frequencies of the HOMs are below the cutoff frequency of the beampipe. For HOMs above the cutoff, the fields may cover regions more than a single cavity, and therefore calculations involving multiple cavities are required to give a correct account of wakefield effects of these modes. To determine traped modes and their effects on the wakefield, we modeled the ILC cryomodule, which consists of eight TTF cavities.

In the frequency domain, we used Omega3P to calculate the trapped modes in the cryomodule for the third dipole band, for mode frequencies above the beampipe cutoff frequency. Sixteen dipole modes with high shunt impedances were evaluated and the calculated damping factors agreed reasonably well with the data measured for cryomodules at DESY. This was the first-ever calculation of trapped modes in the cryomodule. An example dipole mode is shown in Fig. 1. Each mode took about one hour with 1500 processors on the Seaborg computer at NERSC. The calculations were facilitated by advances in parallel linear solvers (see below) so the linear systems resulting from the large computational size of the cryomodule could be solved efficiently on the NERSC machine. A new multi-mode waveguide boundary condition was also implemented in Omega3P to properly terminate the modes at the two ends of the cryomodule. The Maxwell system becomes a nonlinear complex eigenvalue problem, which may be solved using a self-consistent iterative algorithm.

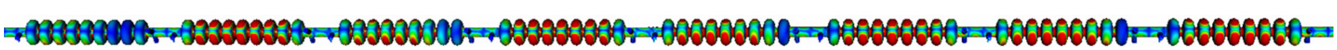

Figure 10: A mode in the third dipole band in the ILC cryomodule. Note that the field pattern rotates along the cryomodule because of the 3-dimensional effects of the couplers located in the interconnecting beampipes. This may produce unwanted x-y wakefield couplings on the bunch.

In the time domain, we used the finite-element code T3P to drive a Gaussian bunch through the cryomodule. The Fourier transform of a monitored signal at the beampipe between two cavities shows a sharp peak close to the monopole cutoff frequency of the beampipe. Visualization of the wakefield excited by the transit of the bunch was rendered using V3D, the 3D visualization tool developed since SciDAC1. Because of the long aspect ratio of the cryomodule, the animation was displayed with zoom-in regions at coupler regions along the cryomodule so that the coupling of fields to the couplers could be clearly demonstrated. The animation was presented at the Supercomputing 2007 Conference on a 96 Mexapixel display consisted of a 5 by 5 tiling of 30" Apple Cinema displays for a total of 25 displays. To further determine the nature of this peak, we used Omega3P to solve for modes at frequencies around the peak, and a localized mode was found in the coupler region between two cavities (see Fig. 2). The shunt impedance and damping factor of this mode indicate that the heat load it generates on the beampipe wall is relatively small and should not be a concern for the ILC linacs.

The work was presented at the ILC Wake Fest 07 Workshop at SLAC, December 11-13, 2007: Rich Lee, Zenghai Li, Greg Schussman, Ravi Uplenchwar, Liling Xiao, Cho Ng, “Multi-Cavity Trapped Mode Simulation”.

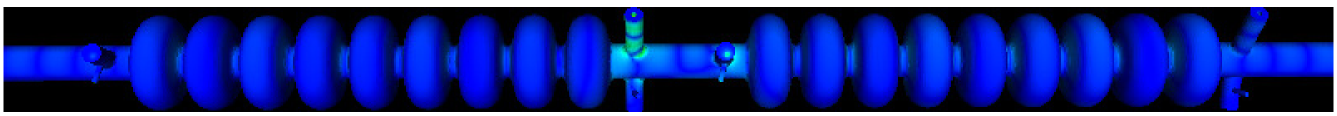

Figure 11: A trapped mode localized in the region between two cavities of the ILC cryomodule.

(2) Multipacting Simulation for Superconducting Cavities (0.5 FTE)

Multipacting is a phenomenon that often limits the performance of superconducting cavities, preventing them from reaching high accelerating gradients. The parallel finite-element tracking code Track3P has been developed to identify the occurrence of multipacting activity in a cavity. The accuracy of Track3P modeling has been improved by the correct treatment of particle localization in curved tetrahedral mesh elements on higher-order surfaces at geometric boundaries. This allows accurate determination of the locations of particle emissions and impacts at a cavity wall, which is crucial to the correct identification of resonant trajectories.

(2.a) The TTF-III power coupler adopted for the ILC baseline cavity design tends to have a long initial high power processing time. The long processing time is believed to result from multipacting in various regions of the coupler. To understand performance limitations during high power processing and identify problematic regions, SLAC has built a flexible high-power coupler test stand to test individual sections of the coupler, including the cold and warm coaxes, the cold and warm bellows, and the cold window. To provide insights for the high power test, detailed numerical simulations of multipacting for these sections were performed using Track3P. Multipacting bands were identified by scanning power levels up to 2 MW. For the cold coax, the simulated multipacting bands agreed very well with measurements (see Fig. 3). Multipacting activities were found in the taper region, but they could be suppressed by applying an axial magnetic field above 300 Gauss. Multipacting activity was not apparent at the bellows.

The work was presented at the ILC Wake Fest 07 Workshop at SLAC, December 11-13, 2007: Lixin Ge, “Multipacting Simulation Using Parallel Code Track3P”.

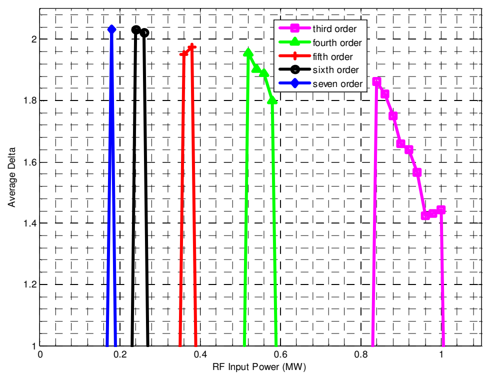

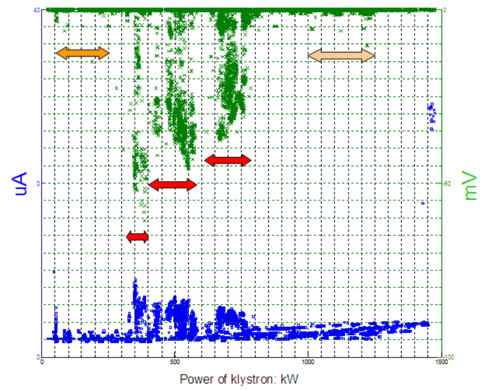

Figure 12: Multipacting activities in the cold coax of the TTF coupler: (Left) Multipacting bands obtained from Track3P; (Right) Electric current measured at a probe during high-power processing.

(2.b) HOM coupler in SNS superconducting cavity

During the commissioning and operation of the SNS superconducting linac, abnormal signals were observed at the HOM coupler. These anomalies are believed to be due to the field emission and multipacting electron loadings in the HOM coupler. To understand these issues, the field enhancement and multipacting were simulated using Omega3P and Track3P for the SNS beta=0.81 cavity. The high field gradient in the upstream HOM coupler could cause field emission and multipacting. Multipacting analysis using Track3P shows two multipacting bands in the gradient range up to 20 MV/m. One of the multipacting bands exists in the gap between the enlarged loop head and the cylindrical side wall of the coupler at field gradient levels from 2.8 to 10 MV/m as shown on the left of Fig. 4. The other MP band exists in the gap between the hook part of the loop and the cylindrical side wall of the coupler as shown on the right of Fig. 4. No MP activities are found in the notch gap. These MP bands are in good agreement with the experimental observations.

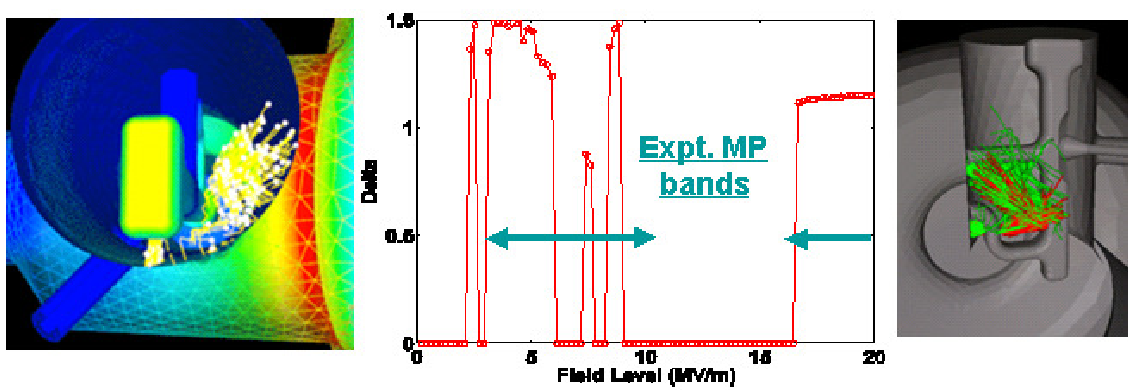

Figure 13. Multipacting activities in the SNS HOM coupler.

(3) Impedance Calculation of the ILC Damping Rngs (0.75 FTE)

The goals of the impedance calculation for the ILC damping rings are to obtain the impedance budget for the vacuum chamber components and to provide a pseudo-Green's function wakefield for beam stability studies. The evaluation of the pseudo-Green's function wakefield necessitates the accurate determination of the short-range wakefield using a bunch length much smaller than the normal bunch length. As the bunch length gets smaller, the computational size becomes larger because the mesh requires refinement to handle the high frequency content of the short bunch.

The conceptual designs of the ILC damping vacuum chamber components were obtained by scaling those in existing machines such as the PEP-II. We used the parallel time-domain code T3P to calculate the pseudo-Green functions of the major vacuum chamber components such as the rf cavity and the beam position monitor (BPM) using a 1 mm bunch (compared with the nominal bunch length of 6 mm). The calculations were carried out on the Bassi computer at NERSC. Fig. 3 shows the field pattern and the pseudo-Green's function wakefield due to the beam transit through the BPM.

It is desirable to obtain the pseudo-Green's function wakefield for even smaller bunch length for more accurate beam stability studies. However, the computational requirements will become prohibitively high. The difficulty can be alleviated by introducing a computational window moving along with the bunch because the short-range wakefield receives contributions only from fields in the longitudinal region covered by the bunch. The two proposed approaches are to have a moving window with p-refinement or a moving window with h-refinement. In the p-refinement method, high-order finite elements are used in the moving window to improve the solution accuracy, and the region outside the window is represented by lowest order finite elements. This drastically reduces the problem size and preliminary studies have shown that CPU time can be reduced by an order of magnitude, compared to using basis functions of uniform order. In the h-refinement method, the moving window has a much more refined grid than the region outside the window. Collaboration with RPI is in progress for moving window with h-refinement to obtain high-quality meshes for computation.

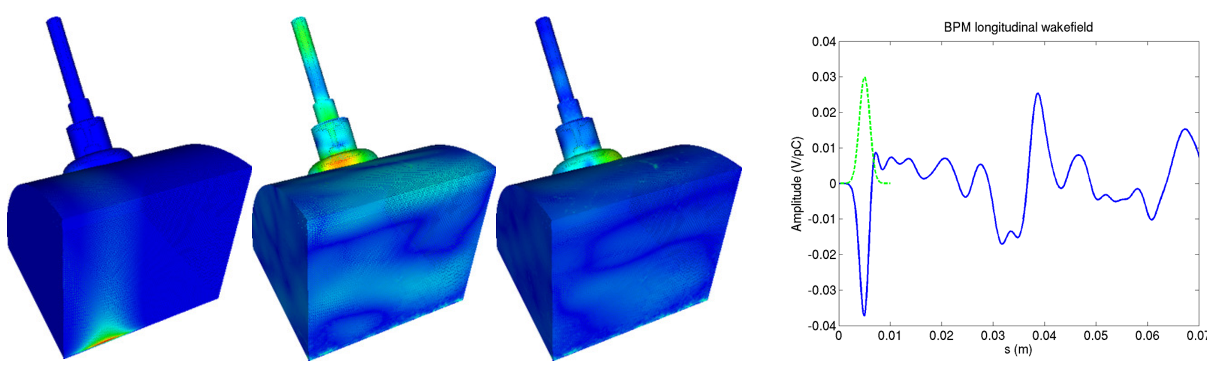

Figure 14: (Left) Transit of a Gaussian bunch through the BPM; (Right) Pseudo-Green's function wakefield for a 1 mm RMS bunch.

Part of this work was presented at the ILC Damping Rings Mini-Workshop at KEK, December 18-20, 2007: Cho Ng, “Status and Plans for the Impedance Calculations of the ILC Damping Rings”.

(4) Development of Effective Domain-Specific Finite-Element Based Multilevel Preconditioners (0.5 FTE)

A hierarchical preconditioner has been developed for solving complex shifted linear systems in nonlinear eigenvalue problems which arise from modeling waveguide-loaded multi-cavity structures. In this preconditioner, the part of the matrix corresponding to the linear basis functions representing a finite-element mesh is factorized using a sparse direct solver. SSOR is then used to combine the linear part and the higher-order part of the matrix in the preconditioning solve. Variations of this hierarchical preconditioner include using real numbers instead of complex numbers, and using single-precision floating point instead of double-precision floating point numbers in the matrix. The method drastically reduces the memory usage compared with the direct factorization of the full matrix so that large problems which could not be solved before on supercomputers are now tractable. The method is suitable for solving large indefinite linear systems emerging from modeling multi-cavity electromagnetic systems such as the ILC cryomodule described in a previous section. The method's scalability is under investigation for possible further parallel performance improvement.

(5) Development of Optimization Software in Shape Determination and Optimization for Unloaded Cavities (0.5 FTE)

In collaboration with TOPS scientists, a shape determination algorithm has been developed to solve for the unknown shape deviations of a cavity from its ideal shape using measured cavity mode frequencies. In the algorithm, the objective is to match the mode frequencies of the deformed cavity model to experimental data through least-squares minimization. The inversion variables are unknown shape deformation parameters that describe perturbations to the ideal cavity shape. The constraint is the Maxwell eigenvalue problem. This nonlinear optimization problem is solved using a line-search based reduced space Gauss–Newton method. The algorithm has been implemented and applied effectively to the ILC cavity data measured at DESY.

The work was published in Journal of Computational Physics, Volume 227 (3), 2008, pages 1722-1738: Volkan Akcelik, Kwok Ko, Lie-Quan Lee, Zenghai Li, Cho Ng, Liling Xiao, “Maxwell Eigenvalue Determination for Deformed Cavities”.

(6) Beam-Beam Simulation for LHC (0.25 FTE)

We have developed a new method to model intrabeam scattering in beam-beam simulations by introducing averaged noise sources and damping. The method is based on the observation that the growth equation for the moments of the particle distribution can be decomposed into a friction part and a noise part, obeying the fluctuation-dissipation theorem. One can thus find a set of coefficients for ring-local Langevin equations for the macroparticles, leading to growth rates in agreement with the Bjorken-Mtingwa (BM) theory.

We have implemented the simplest form of such a Langevin treatment in the Plibb/Lifetime code. It consists of a transformation to the rest frame, the calculation of the growth factors, an Euler step, and a transformation back to accelerator coordinates. It has been verified that the method respects the conserved quantities of BM theory and gives the correct growth rates. The efficiency of the method will be improved by semi-analytic representations of local growth rates, distributing the Langevin kick around the ring, and parallel implementation.

4.8 TECH-X

4.8.1 Progress and Results

This section is divided into three distinct topics: Electromagnetics (EM), Advanced Accelerators (AA) and Electron Cooling (EC). The EM and AA activities are funded by the DOE/HEP office. The EC activity, which is in support of both DOE/HEP and DOE/NP priorities, is funded by the DOE/HEP office in Year 1 (through Aug. 31, 2008), but then will be funded by the DOE/NP office for the remainder of the project. Within each technical topic, the discussion is divided into four subtopics: progress and results, scope modifications, level of effort for supported personnel, and, finally, publications and selected presentations.

4.8.1.1 Electromagnetics (EM)

WBS 1.1.1. Addition of uniformly stable electromagnetic update to VORPAL.
This is still underway and remains part of the planned activities.

WBS 1.1.2. Development of postprocessing software for eigenmode and frequency extraction from time-domain simulations.

This work has been completed and shown (see next section) to accurately extract frequencies even when modes are nearly degenerate. The description of the algorithm has now been accepted for publication [1.4.3]. This work was synergistic with work funded at the University of Colorado.

WBS 1.1.3. Accurate (parts in 10^5) frequency computations of crab cavities.

These calculations have been completed to obtain the crab cavity frequencies as a function of indentation. The comparison results are shown in the figure below.

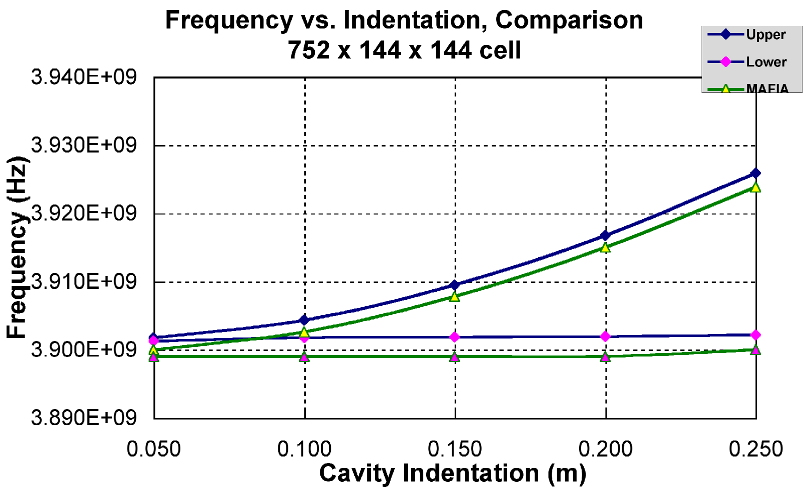

Figure 15. VORPAL and MAFIA crab cavity simulations comparison.

The VORPAL results are in disagreement with previous MAFIA calculations by 3 MHz (out of 3.9 GHz), but the mode separation calculations do agree with previous results. The disagreement cannot be settled by comparisons with experiment, because the cavities are not fabricated with this accuracy in the final frequency. (According to Leo Bellantoni, FNAL crab cavity project leader, cavity frequencies can vary within a range of 10 MHz.) Presentation of this result in the crab cavity collaboration led to a reassessment of previous calculations.

To address this discrepancy, we are approaching this problem from two directions. The group at Lancaster is working on new MAFIA calculations on the crab design to see if higher resolution will give agreement with the VORPAL results. At the same time, the Tech-X group is doing calculations on the A15 cavity, examples of which were fabricated by a more precise machining process, as opposed to the stamping process used for the current NiSn crab cavities.

If the latter approach shows agreement between the VORPAL computations and the experimental measurements, then we will have developed an EM computational standard for cavity frequency computations. This will allow validation of EM codes in non-analytically solvable problems.

4.8.1.2 Advanced Accelerators (AA)

WBS 2.1.1. Benchmarking VORPAL, OSIRIS and QuickPIC

In collaboration with COMPASS team members at LBL and UCLA, we have successfully completed initial 3D benchmarking simulations, considering a characteristic Ti-Sapphire laser pulse entering a constant-density electron plasma and propagating for a short distance. We compared lineouts of the laser electric field, the longitudinal plasma wakefield, and the electron density. Agreement between VORPAL and OSIRIS was found to be quite good, lending confidence in the validity of both codes.

These are the chosen simulation parameters:

electron density: 1.38e+19 cm^-3
plasma frequency: 2.1 e+14 rad/s
laser wavelength: 0.8 microns
laser pulse length: 30. fs (15. fs FWHM)
laser pulse width: 8.2 microns
transverse box: 81.52 microns (along both dimensions)
longitudinal box: 20.5 microns
transverse cells: 0.16 microns
longitudinal cells: 0.04 microns (20 cells per wavelength)
time step: 0.1 fs
of time steps: 1600
mesh size: 512 x 512 x 512
particles per cell: 8
total cells: 1.34e+08
total particles: 1.07e+09

The simulation was run for laser normalized vector potentials of a0 = 0.5, 1, 2 and 4. The transverse profile of the pulse is Gaussian in shape, and the longitudinal profile of the pulse is a fifth-order polynomial, where $f(t) = 10*(t/t\_rise)^3 - 15*(t/t\_rise)^4 + 6*(t/t\_rise)^5$ describes the rise of the pulse, and a symmetric form describes the fall. We used a rise time t_rise=15 fs, which is also equal to the FWHM of the total pulse.

Since OSIRIS initializes the pulse on the grid with the moving window always on, the VORPAL simulation was designed to emit the pulse from the left side of an empty simulation box using a space- and time-dependent boundary condition. Once the pulse has propagated across the empty simulation box to within 5 cells of the right edge, the moving window was turned on. Once the moving window turns on, plasma sweeps into the domain from the right edge of the simulation box. The simulation then runs for 1600 time steps.

VORPAL simulation results are shown in the following images:

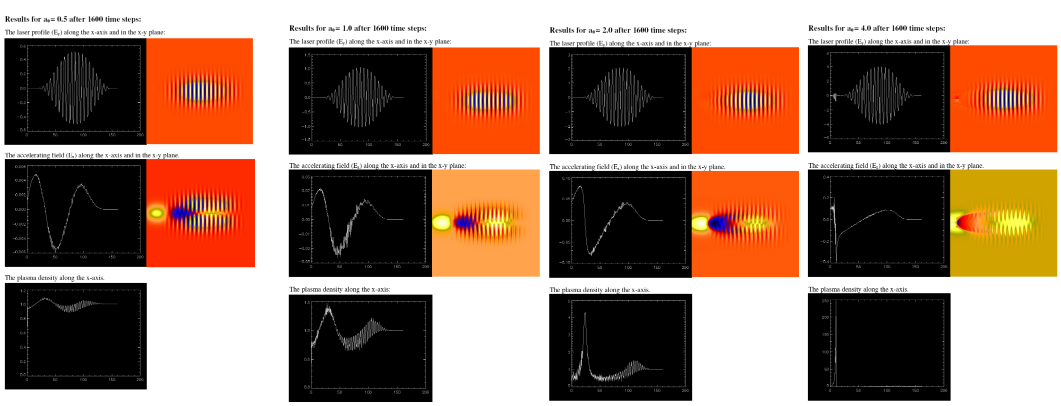

*Figure 16 Normalized laser amplitudes of a0 = 0.5, 1, 2 and 4, left to right, respectively. The electric field amplitudes have been normalized to E0=m_e*c*omega0/e, the electron density has been normalized to the background value, and positions have been normalized to the laser wavenumber k0=2*pi/lambda0.*

Direct comparison of the lineouts from the above VORPAL results with the corresponding OSIRIS simulations show very good agreement. The small differences will be explored in future work. We don't show these comparison plots here. This benchmarking exercise included comparisons with the 3D quasi-static code QuickPIC, which are not discussed here. Future efforts will include the ponderomotive guiding center (PGC) or "envelope" model in VORPAL. In preparation for this work, we have been improving and refactoring the envelope model implementation within VORPAL.

The laser envelope model allows modeling of the laser field in LWFA simulations by the field amplitude envelope rather than by the fields themselves. Since this eliminates the need to resolve the laser wavelength, orders of magnitude speedup are obtained. Proof-of-principle demonstrations of the validity of the model have been completed. The development of the laser envelope model as part of COMPASS involved tighter integration with the VORPAL code base. This development will allow the use of other VORPAL features with the envelope model, including high-order particles, as well as improved diagnostic output.

2.1.2. Demonstrate efficient use of leadership class facilities with VORPAL

VORPAL has been ported to Franklin at NERSC and is also running on ATLAS (a similar opteron-based cluster with many thousands of processors, located at Lawrence Livermore National Lab). For 3D simulations, VORPAL is routinely run on 2,000 processors. For example, to simulate 100 MeV electron acceleration over 3 mm, with experimental parameters relevant to the LOASIS Program at LBL, we find the following (approximate values):

100,000 processor-hours
2,000 processors
100 processor-hours for each data dump
180 s (wall time for each dump)
50 GB of field and particle data per dump
278 MB/s (rate of data streaming in parallel to disk)
40 dumps per run
2 TB of total data
4% overhead for file I/O

Both ATLAS and Franklin use the Lustre parallel file system. We learned that the Lustre file system doesn't work well with the parallel I/O features of HDF5, which are used successfully by VORPAL on several other file systems. This difficulty has been discussed with consultants at both LLNL and NERSC. The work-around in VORPAL was to implement the capability for each MPI process to dump data in serial to separate HDF5 files, then to merge the files in post-processing.

The PGC or envelope model in VORPAL requires use of the Trilinos library from Sandia National Lab (SNL). We haven't yet built VORPAL with Trilinos on Franklin, but this has been done on the NERSC clusters Bassi and Jacquard, where envelope simulations are being tested.

2.2.3. LWFA injection mechanisms

Synergistic with this proposal has been our work investigating new injection mechanisms for LWFA. Multiple publications [2.4.4, 2.4.5, 2.4.6] on this work in collaboration with researchers at the Argonne National Lab and so shows that we are sharing our codes throughout the DOE community.

4.8.1.3 Electron Cooling (EC)

3.1.1. Explore the use of electrostatic PIC in VORPAL for electron cooling simulations

VORPAL uses the Trilinos library from SNL to solve Poission's equation for electrostatic particle-in-cell (PIC) simulations. The latest version of VORPAL has been built with Trilinos on Bassi and Jacquard at NERSC, where we have begun testing the use of PIC for electron cooling simulations. As expected, the minimum impact parameter is limited by the grid size, so the simulated dynamical friction force on an ion is artificially small, as compared to the more accurate simulations with the binary collision model in VORPAL. Even for the binary collision model, electrostatic PIC is used via operator splitting to capture the self-fields of the electron distribution.

3.1.2. Assist SciDAC collaborators in use of VORPAL for cooling simulations

We installed VORPAL on the BNL linux cluster, along with the IDL-based visualization utility VorpalView. We also conducted a VORPAL training session at BNL for V. Litvinenko, A. Fedotov, E. Pozdeyev and others.

We began some simulations of the modulator section for the "coherent" electron cooling (CEC) concept, which combines features of a standard electron cooler and stochastic cooling. The CEC modulator is similar to a standard electron cooler, but rather than imparting a friction force to the ions, the main purpose is to capture the electron density response due to the ions (i.e. the ion wakes), so that these signals can be amplified and used downstream in the "kicker". In collaboration with the COMPASS team at BNL, and with Y. Derbenev of Jefferson Lab, we defined the four key dimensionless parameters that will determine the electron density wake driven by a single ion.

Although physical parameters are not crucial, we decided to use the following as a starting point:

Z_ion = 79 (i.e. fully ionized Au)
n_e = 1.6e16 m^-3 (background electron density)
v_x_rms = 10^6 m/s (rms perpendicular e- velocities)
v_z_rms = v_x_rms/R (rms longitudinal e- velocities)
R ~ 3

The initial simulations, the numerical grid will scale like the following. We use "s" to denote the longitudinal Debye length, while R*s is the transverse Debye length. The numbers 5 and 20 are only rough estimates -- we will make them as small as possible, but sufficiently large that the simulation results are independent of them.

s/dz = (R*s)/dy = (R*s)/dx ~ 5
Lz/s = Lx/(R*s) = Ly/(R*s) ~ 20
Nx=Ny=Nz ~ 100
N_cells ~ 10^6

For electrostatic PIC (and also for particle binning to visualize the results), one requires ~10 particles per cell. This implies ~1e+7 particles. We will have ~125 cells in a Debye sphere, which implies ~1,250 computational particles in a Debye sphere.

An example of the electron density wake generated by an ion can be seen in the figure below:

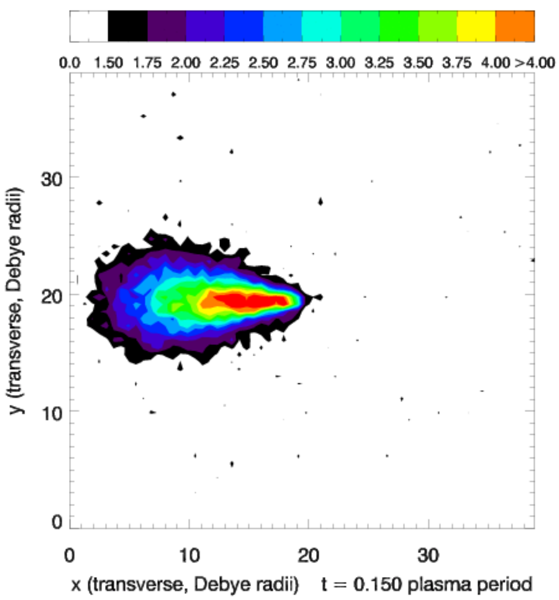

Figure 17. Electron density wake. Positions are normalized to the longitudinal Debye radius. The color contours show the electron density enhancement near (and trailing behind) the ion, which is located at the center of the domain.

4.8.1.4 Year One Scope Modifications

Following tasks were descoped:

EM- Given that the funding of this project is less than originally anticipated, we had to scale back on some activities. With the cutbacks in ILC funding, it was decided to not pursue the two ILC related tasks below.

WBS 1.1.4. Computation of beam energy loss due to wake fields in TESLA cavities.

WBS 1.1.5. Computations of wake fields and momentum kicks associated with end groups.

If we return to these tasks, we now see that we will need first to implement a perfect dispersion algorithm and a moving window for electromagnetic structures first. Hence, these would be new tasks to undertake at the beginning of next year.

There has also been some synergistic work in developing software for visualizing these cavities. At their core has been the development of a visualization schema - a set of attributes written to self-describing data files that allow the contained data to be interpreted and visualized. This allows data for, e.g., cavities to visualized using the VisIT tool being developed under SciDAC. A sample visualization is shown in crabcavwithmesh17mar08.png.

AA - Develop absorbing BCs for laser envelope in ponderomotive guiding center (PGC) PIC model.
EC - Generalize MD algorithm to use electrostatic PIC for long-range interactions.

4.9 TJNAF

4.9.1 Progress and Results

4.9.1.1 Electromagnetics (EM)

Budget: The funded budget for EM simulations of Year One task is $40K including overhead. So far in the first six months, we haven't started using this money yet. The reason is that we are hiring a Post Doc to do the ComPASS work. This budget only covers half of this FTE. The hired Post Doc will be in position on April 14, 2008

Resources Development: Upon we received the SciDAC fund, we identified that the scope of the ComPASS work and other demand from different projects need a full time Post Doc position on the SRF structure. We immediately started hiring process and organized an interview panel including the members from SRF, CASA, FEL and Engineering departments. After reviewing 14 applications, three candidates were called up for interviews. The final candidate was determined by the highest total score from all panel members. We fortunately completed this process before the Thanksgiving holidays. Before the offering letter sending out last December, we met hiring freeze due to the congress's continue resolution. After review our funded budget situation in SciDAC and other projects, we decided that this position would be less impact to the nuclear physics budget. So a strong petition to our senior managers had been made and finally had this position exempt from this hiring freeze.

Before the Post Doc coming to the post, we have installed the Tech-X's VORPAL and SLAC's Omega3P and S3P software on our local work station (stand along node, Linux OS). Additional library, graphic, mesh generator and post possessing tool software also have been compiled and installed properly. Some test runs were trialed. We invited Dr. Chet Nieter from Tech-X to give a tutorials session and talk related to the VORPAL code and its development. Separately, we also invited Dr. Zenghai Li to visit JLab and gave tutorials and seminar on the SLAC's code family development.

As one partner of ComPASS project, JLab does not do the code development but code benchmark. We have provided equally to our code development partners our design, prototype and measurement data from our R&D and production projects, as well as experiment data from our superconducting accelerators during the course of collaboration.

All these activities carried out were funded by other sources than SciDAC.

Dipole Passband HOM damping simulation on fully dressed JLab High Gradient cavity (See Figure 1) has been done by Zenghai Li, Liling Xiao and Akcelik Volkan at SLAC on NERSC computer.

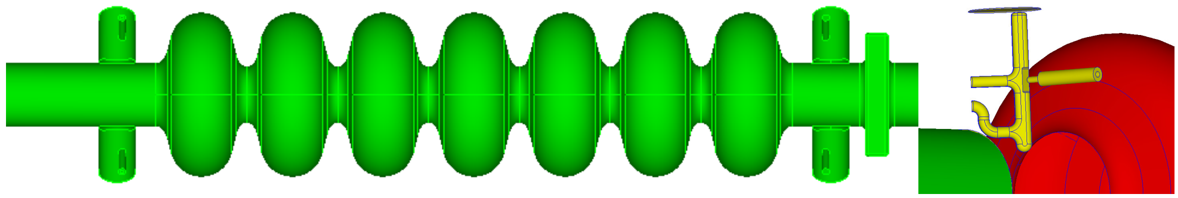

Figure 18: Omega3P/S3P simulated JLab 7-cell High Gradient SRF cavity. Left: fully dresses cavity. Right: detail of HOM coupler structure.

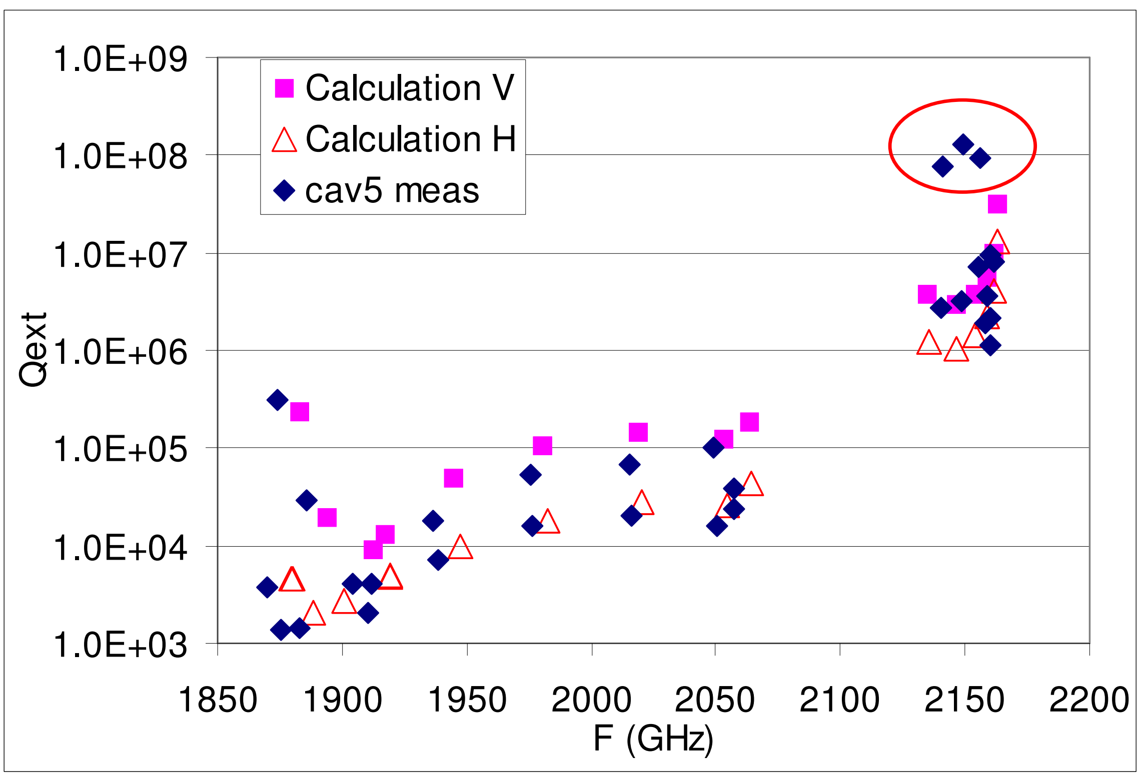

Figure 19: Omega3P simulated external Q of first two dipole passband modes, their polarizations and comparison with measured data on the cold cavity. The red circle indicates three high-Q modes. Two of them caused BBU at CEBAF.

The simulated HOM coupling (Qext) needs great detail of structure (large amount of meshes) and high frequency mode information (a lot of CPU time). The agreement from the simulation to the measurement is great except three trouble modes indicating a problem either due to the couplers or to the field profile titled to the FPC side.

We further investigated the cavity's cell imperfection verses external Q change. Table 1 is an initial investigation result. Due to a strong cell-to-cell coupling on these dipole modes, the field flatness and external Q changes are small.

Table1: JLab HG cavity cell imperfection verses HOM external Q.

	5PI/7		4PI/7	
	F (GHz)	Qext (10^6)	F (GHz)	Qext (10^6)
Design (pickup gap=0.75mm)	2.146	2.91	2.154	3.74
1st iris (-1mm)	2.147	4.07	2.155	5.51
Mid-cell length (-1mm)	2.146	2.85	2.155	4.10
3rd-cell length (-1mm)	2.148	2.16	2.1546	3.97

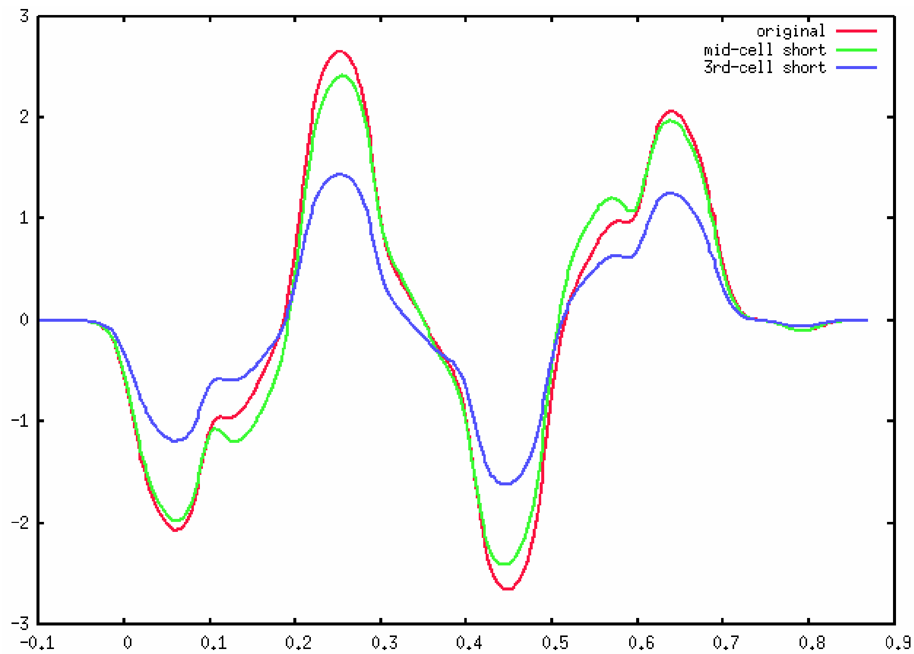

Figure 20: Electric field profile change on the 4pi/7 mode of TM110 passband due to the deformation in Table 1.

A TDR measurement on the cavity#5's HOM cables indicates an abnormal signal. This leads the simulation investigation on the different HOM coupler's pickup probe lengths.

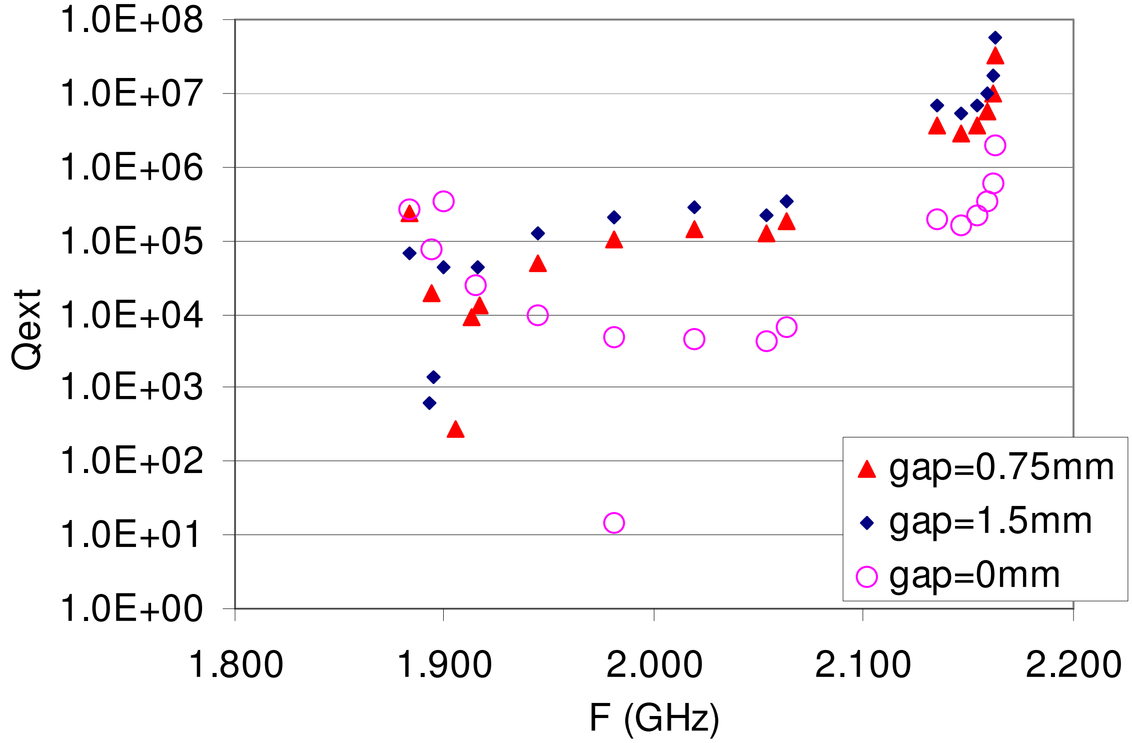

Figure 21: Omega3P simulation indicates a larger pickup gap will cause an overall Qext increase on all dipole passband modes.

This is still an on-going investigation, no cause of this HOM damping problem has been brought into conclusion yet at this report time. This collaboration will be enhanced further in the ComPASS project work frame.

Other work in progress area is on the JLab's Year Three work plan in ComPASS: 3D multipactoring simulation in a SRF cavity. We were able to provide Tech-X's partners the multipactoring data on a TE011 coaxial type cavity. This worked is funded by Tech-X's SBRI phase II proposal (DOE SBIR grant DE-FG02-05ER84172) supported by JLab. We have provided Tech-X with cavity geometry, cold experiment and secondary electron yield data as well as the initial result analysis. The RF field inside of cavity was first excited by a current source in time domain. A space-charge self-consistent particle tracking for electrons were tracked. The RF field amplitude, phase and primary electron launching position were scanned. A multiple secondary electron trajectories can be found. Figure 6's result is one of such trajectories generated by Selma Cetiner, Chet Nieter and Peter Stoltz of Tech-X Corporation.

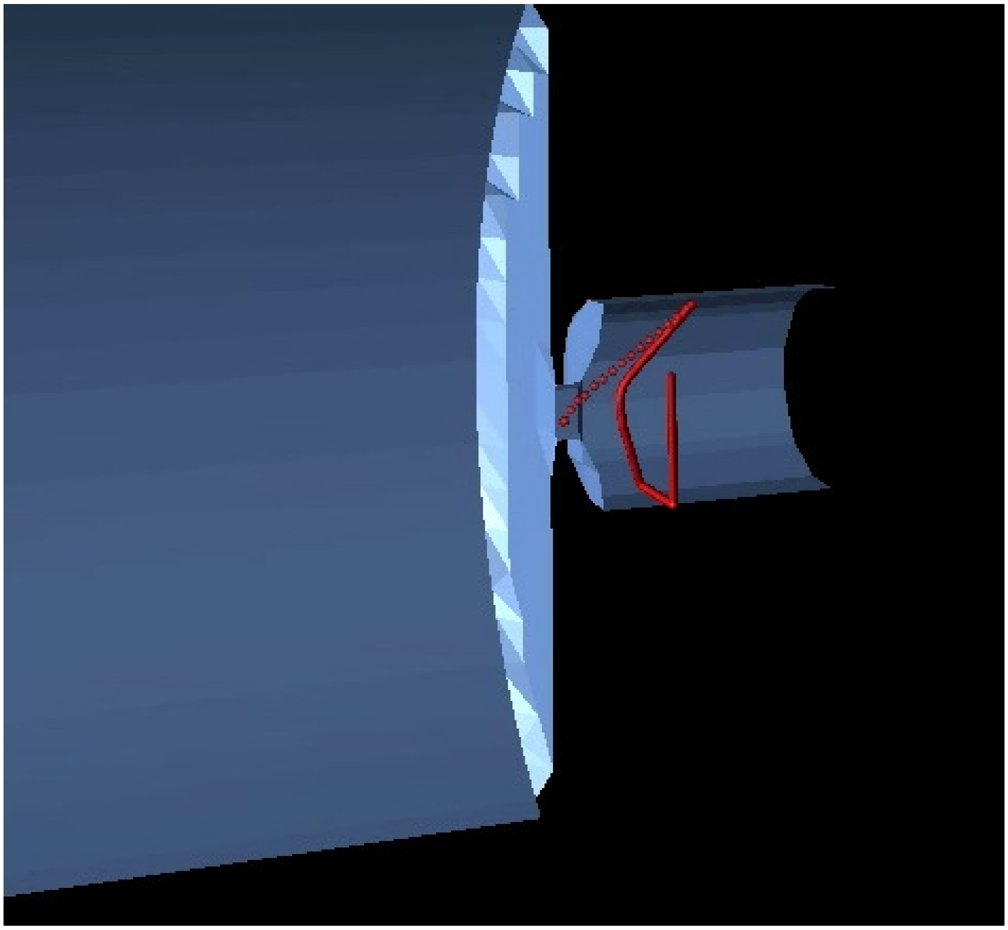

Figure 22: Multipactoring trajectory in the coaxial coupler of a JLab TE011 coaxial type cavity simulated by Tech-X's VORPAL code.

5. Work Plan in Next 6 Months:

Once the Post doc is on board, we are going to start following:

- Complete the SLAC's code family installation on our local work station including T3P, Track3P, Gun3P, Pic3P and TEM3P.
- Send him to SLAC ACD group to learn these codes for one week.
- Start benchmarking wake fields and RF-thermal problem in a SRF cavity with SLAC's and Tech-X's codes.
- Work with our CASA members on the 64-CPU local cluster computer initiative.

4.9.1.2 Year One Scope Modifications - EM

On the contrary, started from last November, we met beam breakup (BBU) problem occasionally depending on the particular optical setup. The threshold injector current varied from 38 to 100 µA up to 140 µA without BBU. The further beam based BBU test and RF measurements confirmed that this BBU was caused by either one of two high-Q HOMs inside a single High Gradient (HG) cavity (#5) in Renascence cryomodule. Comparing to other 7 cavities within the cryomodule, this cavity exhibits an abnormal HOM damping performance than others and being higher than the impedance budget on two dipole passband modes. The QA process to eliminate this problem was bypassed due to a schedule constraint. To understand this problem, we decided to change the ComPASS priority in Year One task with this one: Identifying HOM modes and coupling issues associated with cavity tunings, trapped modes, or misshaped coupler geometries. Confirming the data obtained from the Low Loss cavities built with different HOM coupler configuration have sufficient HOM damping for the 12GeV beam current.

The major tools to simulate such problem and partially critical to our analysis are SLAC's Omega3p and S3P codes. The complementary model and experimental data have also been provided and guided to this simulation work. The simulations are being carried out at SLAC, ACD group.

4.9.1.3 Beam Dynamics (BD)

We have set up a BeamBeam3D simulation deck for ELIC with a single IP and head-on collision and performed simulations with nominal design parameters. Extensive work was focused on carrying out numerical convergence tests, so as to reach reliable simulation results which are convergent against simulation parameters, such as longitudinal slice numbers, transverse mesh size, and number of macro-particles. We then conducted thorough parameter dependence studies of ELIC luminosity and beam-beam instabilities. We identified parameter regimes for the electron and proton beam currents where nonlinear beam-beam interaction and coherent instability becomes important. We further scanned the betatron tune map for searching good working point. We are currently extending our effort to multiple IP simulations with an ELIC deck accommodating up to 4 IPs on a figure-8 ring and plan to conduct multi-bunch multi-IP simulation studies in the second half of this year.

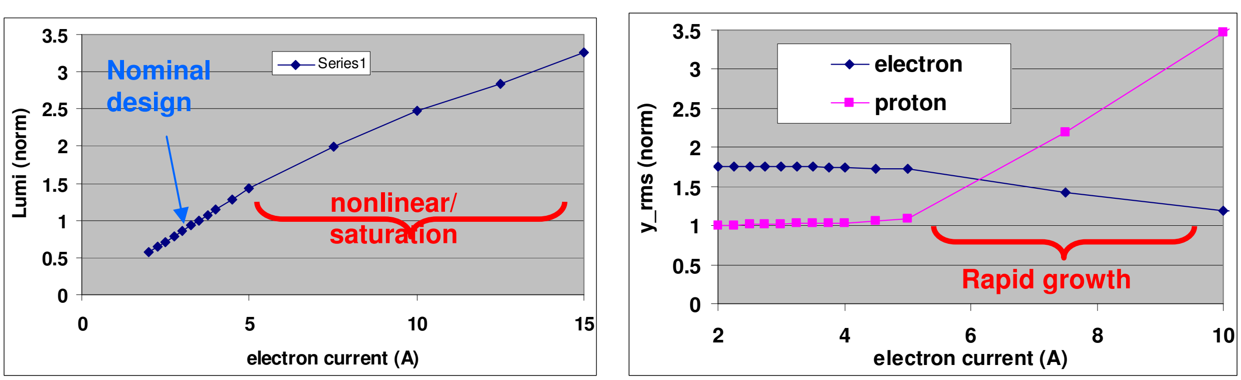

Figure 23: (a) ELIC luminosity (normalized to the design value, 7.8x10^{34}) as a function of electron current, and (b) the growth of vertical size of proton bunch (normalized by the design value)

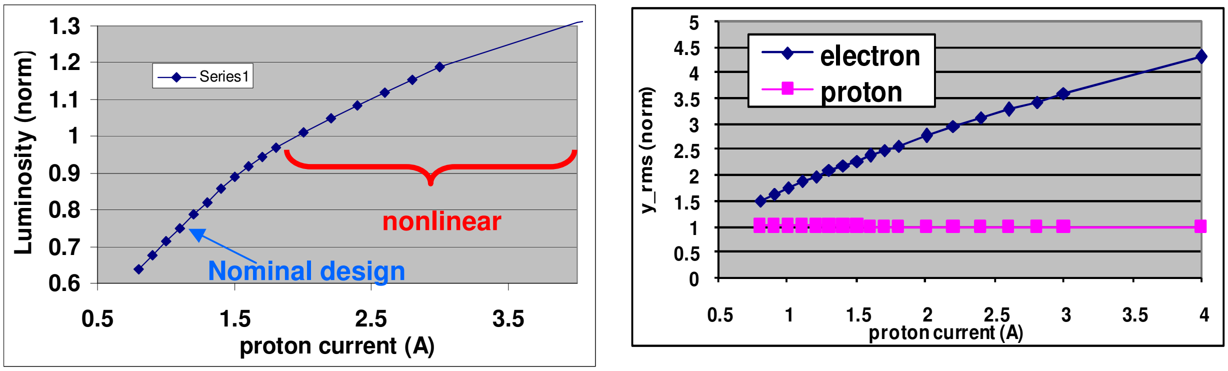

Figure 24: (a) ELIC luminosity (normalized to the design value, $7.8x10^{34}$) as a function of proton beam current, and (b) the growth of electron bunch vertical size (normalized by the design value)

Electron cooling

In the area of ELIC electron cooling modeling, we are investigating applicability of the existing code suites to ERL circulator-ring based ELIC electron cooling design. After this pre-simulation investigation we will move on to the cooling simulations with VORPAL and BetaCool code suites in the second half of the present funding year.

4.9.2 Year One Scope Modifications

EM - Original JLab's Year One work plan according to the reduced budget in the ComPASS proposal was "Provide Low Loss cavity shape for the 12 GeV Upgrade for benchmark VORPAL, Omega-3P, T3P on its simulated modes, short range wake fields, and heat generation due to the rf surface resistance by comparing to analytical formulae, estimates, experimental data".

This priority was developed during the urgent need of CEBAF 12 GeV Upgrade R&D project. Heating and thermal runaway problem on coaxial HOM coupler pickup probes has been encountered in the SRF cavities of "Renascence" cryomodule. At that time, no simulation tool including commercial (ANSYS) code is available to understand this problem. Instead, we used engineering solution to remove the HOM couplers' probes on the Fundamental Power Coupler (FPC) side and modified the feedthroughs of HOM couplers on the Field Probe (FP) side and furthermore added the thermal interception of HOM cables inside of cryomodule. By doing that we were able to achieve the design accelerator gradient in 20MV/m.

4.10 UCLA

4.10.1 Progress and Results

Pipelining

The ambitious work statement for the UCLA effort duing the first year was:

1. Finish adding PML and splines into OSIRIS
2. Isolate bottlenecks in QuickPIC loop.
3. Finish basic version of pipelining of QuickPIC.
4. Validate OSIRIS, Vorpal, and QuickPIC against each other.
5. Model experiments.
6. Study beam loading.

During the first 6 months the effort is progressing extremely well. Below we summarize some of the progress. Some of the work below is not directly funded by SciDAC.

Pipelining and improvements to QuickPIC:

In the QuickPIC computation cycle, the plasma response to the driver is evaluated by sweeping a slice of plasma along the longitudinal direction and calculating the trajectories of the plasma particles as the slice moves from the front of the driver to the back. This setup allows a software pipelining technique to be used, which can dramatically decrease the turn-around time of a PWFA/LWFA or electron cloud simulation by adding more processor groups to the pipeline. We have implemented the pipelining algorithm into both the basic and full quasi-static version of QuickPIC for the particle beam driver. The HDF diagnostic routine has also been modified to allow merging data of each processor group on the fly without post-processing. This pipelining algorithm has been successfully tested on NERSC platforms with 2048 processors in which a maximum of 64 processor groups was used. The pipelining algorithm is verified to produce the same result as the non-pipelining version. Performance measurement shows that the speedup of the 2D plasma slice solver in the pipelining mode is nearly ideal, this is because the data transfer between two successive groups is relatively inexpensive compared to the time spent on the solver itself and the transfer also overlaps with the computation. When the 3D beam pusher is included, the overall efficiency of each processor group in the pipelining mode reaches 85% relative to the non-pipelining mode, thus using 64 processor groups leads to a 54 times reduction in turn-around time for a long simulation. A description of the pipelining routine will be submitted to J. Comp. Physics shortly. We also identified two key factors for improving the efficiency and processor number scaling of the pipelining algorithm, they are (1) choosing a 2D domain decomposition strategy in the 3D beam pusher and (2) implementing load-balancing in the same pusher. We have also begun to improve the predictor corrector routine for both preionized and self-ionized plasmas.

Beam Loading:

A theory of how electrons can be loaded into the nonlinear wakefield generated by either a laser or particle drive beam has been developed. The maximum amount of charge that can be loaded and the optimum shape are derived and confirmed by OSIRIS simulations. In the ideal case, the efficiency of beam-loading is shown to be 100%. A preliminary plasma wakefield afterburner design based on this theory has been simulated in a QuickPIC simulation on NERSC. In the simulation, the drive electron beam has 4.4×10^{10} electrons and a length of 58 microns, the trailing electron beam has 1.7×10^{10} electrons and a length of 22 microns. The drive beam has a triangular longitudinal current profile while the trailing beam has a inverse triangular profile, this choice of the beam profiles has lead to complete flattening of the decelerating and the accelerating fields in the plasma wake. Therefore, zero energy spread in achieved for the accelerated trailing beam! A paper has been submitted to Phys. Rev. Lett.

Benchmarking codes:

In collabortation with Tech-x and LBNL, we have begun benchmarking of the three major simulation codes used in advance accelerator modeling effort, OSIRIS, VORPAL and QuickPIC.. A series of 3D benchmarks covering the linear and nonlinear regimes of LWFA problem were run with each code. In these benchmarks, the laser pulse has the following fixed parameters, laser wavelength λ=0.8 micron, pulse duration τ = 60 fs, spot size w_0 = 8.2 micron, and the magnitude of the normalized vector potential of the laser are a0=0.5, 1, 2 and 4. The plasma density is 1.38×10^{19} cm^{-3}. The simulation box size is 80×80×20 microns with 512×512×512 grid points. The time step is 0.2075 ω_0^{-1}. The longitudinal wake electric field and the laser electric field in OSIRIS and VORPAL simulation at 800 and 1600 time steps were compared with each other. In addition, longitudinal electric field from QuickPIC simulation is also compared with the full PIC results. The benchmarks show that the OSIRIS and VORPAL results are nearly identical. QuickPIC results are also seen to agree extremely well with the full PIC results for a0=0.5, 1, 2 cases, while for the a0=4 case deviation is seen near the spike of the accelerating field. We will next run benchmarks for which the laser evolves, as well as test ideas for how to improve make QuickPIC agree betterfor the larger values of a_0. QuickPIC and OSIRIS are also being compared to the code WAKE. The quasi-static codes can be more than two orders of magnitude faster than the full PIC codes.

Modeling experiments:

We have conducted several high resolution QuickPIC simulations to study the relevant physics in a plasma afterburner. The design of the beam parameters is based on the beam-loading theory described above. The spot size of the drive beam is 10 microns to avoid significant ion motion, while the trailing beam spot size has to be sub micron for collider application. In the simulation, we use 2048×2048×256 grids with box size of 600×600×270 microns. The transverse resolution is ~0.3 micron which allows us to simulate a 0.6 micron spot size trailing beam. Both the drive beam and the trailing beam have initial energy of 250 GeV. After propagating about 4 meters in the plasma, the trailing beam has gain around 100 GeV energy. We have conducted simulations with/without ion motion and synchrotron radiation loss to investigate their effects on beam acceleration and propagation. We have also assisted in carrying out OSIRIS simulations of laser guiding experiments conducted at UCLA.

Scaling PIC codes to 10,000 processors:

We have done extensive benchmarks with the UPIC Framework on large systems to identify areas where current algorithms can be improved to enable scaling of PIC codes to 100,000 processors. This is directly applicable to the QuickPIC code, which depends on this Framework, but new successful algorithms can also be implemented in the other PIC codes. A billion particle benchmark simulation with a 512x256x512 grid, using an electromagnetic, relativistic plasma model has been run successfully on an Opteron-based cluster with Infiniband. This strong scaling study in which the number of processors was varied from 128 to 8192 while keeping the particle size fixed showed good scaling (92%) for the particle part of the calculation (which typically dominates PIC codes). The FFT part scaled well up to 4096 processors, then saturated at 8192. Top processing speeds of more than10 billion particles per second (for an entire step including field solve etc.) were obtained. To enable scaling to 100,000 processors, two approaches are being tested. First, strategies to improve existing algorithms in the particle manager and the FFT were identified. Second, a mixed MPI/threaded programming model has been implemented and is now being tested.

4.10.2 Year One Scope Modifications

5 Change Log

Revision No.	Description/Pages Affected	Effective Date
Revision 0.0		

6 Project Office Contact

This report is compiled by:

Dr. Bakul Banerjee bakulb@fnal.gov
COMPASS Project
Fermi National Accelerator Laboratory
POB 500, Batavia, IL 60510
Tel: +1 630 840 5251

7 Publications

Papers

Each PI is responsible for reporting publications and presentations made based on the research activities assocated with the project along with the accomplishment reports.

Reported by ANL: [1] , [2] , [3] , [4]. , [5]

Reported by LBNL: [6] , [7]., [8]., [9]., [10]

Reported by Tech-X : **EM -** [11]., [12], [13] , **AA -** [14] ,[15]., [16]., [17] , [18]., [19]., **EC:** [20] , [21]., [22].

Reported by LBNL: [23]), [24]., [25] , [26] , [27]

Presentations

Reported by ANL: [28] , [29]

.

---

[1] J. Cary, P. Spentzouris, J. Amundson, L. McInnes, M. Borland, B. Mustapha, B. Norris, P. Ostroumov, Y. Wang, W. Fischer, A. Fedotov, I. Ben-Zvi, R. Ryne, E. Esarey, C. Geddes, J. Qiang, E. Ng, S. Li, C. Ng, R. Lee, L. Merminga, H. Wang, D. Bruhwiler, D. Dechow, P. Mullowney, P. Messmer, C. Nieter, S. Ovtchinnikov, K. Paul, P. Stoltz, D. Wade-Stein, W. Mori, V. Decyk, C. Huang, W. Lu, M. Tzoufras, F. Tsung, M. Zhou, G. Werner, T. Antonsen and T. Katsouleas, *COMPASS, the Community Petascale Project for Accelerator Science and Simulation, a Broad Computational Accelerator Physics Initiative*, , Journal of Physics: Conference Series 78 (2007).

[2] D.R. Dechow, B. Norris, and J. Amundson, *The Common Component Architecture for Particle Accelerator Simulations*, Proceedings of HPC-GECO/CompFrame'07, October 21-22, 2007, Montreal, Quebec, Canada, ACM, 2007.
[3] "Superconducting accelerating structures for very low velocity ion beams", J. Xu, K. W. Shepard, P. N. Ostroumov, J. D. Fuerst, G. Waldschmidt and I. V. Gonin, Phys. Rev. ST Accel. Beams 11, 032001 (2008).
[4] "A Parallel 3D Poisson Solver for Space Charge Simulation in Cylindrical Coordinates", J. Xu, P.N. Ostroumov and J.A. Nolen, Comp. Phys. Comm. 178 (2008) 290.
[5] "A Parallel Wavelet-Based Poisson Solver for Beam Dynamic Simulations", J.Xu, M. Israeli, P.N. Ostroumov and J.A. Nolen, submitted to SIAM Journal on Scientific Computing.
[6] J.-L. Vay, *Phys. of Plasmas* **14**, 1 (2008)
[7] J.-L. Vay, Phys. Rev. Lett., 98 (2007) 130405
[8] J. P. Boris, Relativistic plasma simulation-optimization of a hybrid code, Proc. Fourth Conf. Num. Sim. Plasmas, Naval Res. Lab., Wash., D. C., 3-67, 2-3 November 1970
[9] C. G. Darwin, Philos. Mag. 39, 537 (1920)
[10] J. D. Jackson, Classical Electrodynamics, (Wiley, New York,1975), 2nd Edition, pp. 593-595.
[11] C. Nieter, J. R. Cary, G. R. Werner, D. N. Smithe, and P. H. Stoltz, "Application of Dey-Mittra conformal boundary algorithm to 3D electromagnetic modelin," J. Comp. Phys., in review (2008)
[12] J.R. Cary, P. Spentzouris, J. Amundson, L. McInnes, M. Borland, B. Mustapha, B. Norris, P. Ostroumov, Y. Wang, W. Fischer, A. Fedotov, I. Ben-Zvi, R. Ryne, E. Esarey, C.G.R. Geddes, J. Qiang, E. Ng, S. Li, C. Ng, R. Lee, L. Merminga, H. Wang, D.L. Bruhwiler, D. Dechow, P. Mullowney, P. Messmer, C. Nieter, S. Ovtchinnikov, K. Paul, P. Stoltz, D. Wade-Stein, W.B. Mori, V. Decyk, C.K. Huang, W. Lu, M. Tzoufras, F. Tsung, M. Zhou, G.R. Werner, T. Antonsen, T. Katsouleas, "COMPASS, the COMmunity Petascale project for Accelerator Science and Simulation,

a broad computational accelerator physics initiative," J. Physics: Conf. Series 78, 012009 (2007). http://www.iop.org/EJ/article/1742-6596/78/1/012009/jpconf7_78_012009.pdf
[13] G. R. Werner and J. R. Cary, "Extracting Degenerate Modes and Frequencies from Time Domain Simulations," J. Comp. Phys., accepted (2008), http://dx.doi.org/10.1016/j.jcp.2008.01.040
[14] C.G.R. Geddes, D.L. Bruhwiler, J.R. Cary, E. Cormier-Michel, E.Esarey, C.B. Schroeder, W.A. Isaacs, N. Stinus, P. Messmer, A. Hakim, K. Nakamura, A.J. Gonsalves, D. Panasenko, G.R. Plateau, Cs. Toth, B. Nagler, J. van Tilborg, T. Cowan, S.M. Hooker and W.P. Leemans, "Laser wakefield simulations towards development of compact particle accelerators," J. Physics: Conf. Series 78, 012021 (2007). http://www.iop.org/EJ/article/1742-6596/78/1/012021/jpconf7_78_012021.pd
[15] K. Paul, D.L. Bruhwiler, J.R. Cary, P. Messmer, P.J. Mullowney, E. Cormier-Michel, E.H. Esarey, C.G.R. Geddes, W.P. Leemans and C. Schroeder, "Benefits of Higher-Order Particle Shapes in the Electromagnetic PIC Code VORPAL," presented at the 49th Annual Meeting of the Division of Plasma Physics (Orlando, Nov. 12, 2007)
[16] D.L. Bruhwiler, J.R. Cary, C.G.R. Geddes, B. Cowan, K. Paul, P.J. Mullowney, P. Messmer, D.A. Dimitrov, E. Cormier-Michel, E.H. Esarey and W.P. Leemans, "Simulation tools for 10 GeV and beyond," presented at the BELLA Proposal Review (Berkeley Lab, Oct. 11, 2007), DOE Office of High Energy Physics, status CD0
[17] Baifei Shen, Yuelin Li, Karoly Nemeth, Hairong Shang, Robert Soliday, Robert Crowell, Edward Frank, William Gropp, and John Cary, "Triggering wave breaking in a laser plasma bubble by a nanowire," Phys. Plasmas V 14, 053115 (2007).
[18] Baifei Shen, Yuelin Li, M. Y. Yu, and John Cary, "Bubble regime for ion acceleration in a laser driven plasma," Phys. Rev. E V. 76, 055402 (2007)
[19] Karoly Nemeth, Baifei Shen, Yuelin Li, Robert Crowell, Katherine C. Harkay, and John R. Cary, "Laser driven coherent betatron oscillation in a laser-wakefield cavity," Phys. Rev.Lett. 100, 095002 (2008)
[20] D.L. Bruhwiler, "Parallel Simulation of Electron Cooling Physics for Relativistic Ion Beams," SciDAC 2007 Workshop (Boston, June, 2007), invited. http://ornlmedia.ornl.gov/scidac07/6-26/1030-1200/David_Bruhwiler.wmv
[21] D.L. Bruhwiler, G.I. Bell, A. Sobol, P. Messmer, J. Qiang, R. Ryne, W. Mori, A. Fedotov, I. Ben-Zvi, V. Litvinenko, S. Derbenev, R. Li, Y. Zhang and L. Merminga, "Parallel Electron Cooling Simulations for SciDAC," presented at Fermilab (Sep. 17, 2007)
[22] D.L. Bruhwiler, G.I. Bell, A. Sobol, V. Litvinenko, I. Ben-Zvi, A. Fedotov and S. Derbenev, "VORPAL Simulations of the Modulator for a Coherent Electron Cooler," presented to BNL Electron Cooling Group (Feb. 20, 2008)
[23] J.-L. Vay, *Phys. of Plasmas* **14**, 1 (2008
[24] J.-L. Vay, Phys. Rev. Lett., 98 (2007) 130405
[25] J. P. Boris, Relativistic plasma simulation-optimization of a hybrid code, Proc. Fourth Conf. Num. Sim. Plasmas, Naval Res. Lab., Wash., D. C., 3-67, 2-3 November 1970.

[26] C. G. Darwin, Philos. Mag. 39, 537 (1920).
[27] J. D. Jackson, Classical Electrodynamics, (Wiley, New York,1975), 2nd Edition, pp. 593-595.
[28] D. Dechow, D. Abell, P. Stoltz, (Tech-X, Boulder, Colorado), L. McInnes, B. Norris (ANL, Argonne, Illinois), J. Amundson (Fermilab, Batavia, Illinois), *A Beam Dynamics Application Based on the Common Component Architecture*, Particle Accelerator Conference (PAC07) Albuquerque, NM, June 25-29, 2007.
[29] D. Dechow, D. Abell, P. Stoltz, (Tech-X, Boulder, Colorado), L. McInnes, B. Norris (ANL, Argonne, Illinois), J. Amundson (Fermilab, Batavia, Illinois), *A Beam Dynamics Application Based on the Common Component Architecture*, Particle Accelerator Conference (PAC07) Albuquerque, NM, June 25-29, 2007.